



Junos[®] OS

Protocol Family and Interface Address Properties

Release

14.1



Published: 2014-05-15

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

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

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

Junos[®] OS Protocol Family and Interface Address Properties

14.1

Copyright © 2014, Juniper Networks, Inc.

All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

	About the Documentation	xxiii
	Documentation and Release Notes	xxiii
	Supported Platforms	xxiii
	Using the Examples in This Manual	xxiv
	Merging a Full Example	xxiv
	Merging a Snippet	xxv
	Documentation Conventions	xxv
	Documentation Feedback	xxvii
	Requesting Technical Support	xxviii
	Self-Help Online Tools and Resources	xxviii
	Opening a Case with JTAC	xxviii
Part 1	Overview	
Chapter 1	Protocol Family and Interface Address Properties	3
	Protocol Family Configuration and Interface Address Statements	3
Part 2	Configuration	
Chapter 2	Protocol Family and Interface Address Properties	9
	Setting the Protocol MTU	9
	Disabling the Removal of Address and Control Bytes	10
	Disabling the Transmission of Redirect Messages on an Interface	10
	Applying Policers	11
	Applying Aggregate Policers	12
	Example: Applying Aggregate Policers	13
	Applying Hierarchical Policers on Enhanced Intelligent Queuing PICs	14
	Hierarchical Policer Overview	15
	Hierarchical Policing Characteristics	16
	Configuring Hierarchical Policers	16
	Configuring a Single-Rate Two-Color Policer	17
	Configuring a Single-Rate Tricolor Policer	18
	Configuring a Two-Rate Tricolor Marker Policer	19
	Applying a Filter to an Interface	20
	Defining Interface Groups in Firewall Filters	22
	Filter-Based Forwarding on the Output Interface	22
	Example: Applying a Filter to an Interface	23
	Enabling Source Class and Destination Class Usage	25
	Examples: Enabling Source Class and Destination Class Usage	29

Configuring the Protocol Family	33
IPv6 Overview	35
IPv4-to-IPv6 Transition	35
VRRP Properties	35
Configuring the Interface Address	36
Configuring Interface IPv4 Addresses	37
Operational Behavior of Interfaces when the Same IPv4 Address is Assigned to Them	38
Configuring Interface IPv6 Addresses	39
Configuring ICCP for MC-LAG	40
Configuring IPCP Options	42
Configuring an IP Address for an Interface	42
Negotiating an IP Address Assignment from the Remote End	43
Configuring an Interface to Be Unnumbered	43
Assigning a Destination Profile to the Remote End	43
Configuring an Unnumbered Interface	44
Configuring an Unnumbered Point-to-Point Interface	44
Example: Configuring an Unnumbered Point-to-Point Interface	45
Configuring an Unnumbered Ethernet or Demux Interface	45
Configuring a Preferred Source Address for Unnumbered Ethernet or Demux Interfaces	46
Configuring Static Routes on Unnumbered Ethernet Interfaces	47
Restrictions for Configuring Unnumbered Ethernet Interfaces	48
Example: Configuring an Unnumbered Ethernet Interface	48
Example: Configuring the Preferred Source Address for an Unnumbered Ethernet Interface	49
Example: Configuring an Unnumbered Ethernet Interface as the Next Hop for a Static Route	49
Configuring Default, Primary, and Preferred Addresses and Interfaces	50
Configuring the Primary Interface for the Router	50
Configuring the Primary Address for an Interface	51
Configuring the Preferred Address for an Interface	51
Configuring Unicast RPF	52
Unicast RPF and Default Routes	53
Unicast RPF Behavior with a Default Route	53
Unicast RPF Behavior Without a Default Route	54
Unicast RPF with Routing Asymmetry	54
Configuring Unicast RPF Strict Mode	54
Example: Configuring Unicast RPF Strict Mode	55
Configuring Unicast RPF Loose Mode	56
Example: Configuring Unicast RPF Loose Mode	57
Configuring Unicast RPF Loose Mode with Ability to Discard Packets	58
Example: Configuring Unicast RPF Loose Mode with Ability to Discard Packets	58
Configuring Unicast RPF on a VPN	59
Example: Configuring Unicast RPF on a VPN	60

	Example: Configuring Unicast RPF	60
	Configuring Targeted Broadcast	61
	Configuring Targeted Broadcast and Its Options	61
	Display Targeted Broadcast Configuration Options	62
	Forward IP Packets On the Egress Interface and To the Routing Engine	62
	Forward IP Packets On the Egress Interface Only	62
	Understanding Targeted Broadcast	63
	Example: Configuring Unicast Reverse-Path-Forwarding Check	64
	Understanding Unicast Reverse Path Forwarding	64
	Example: Configuring Unicast Reverse-Path-Forwarding Check	65
Chapter 3	Network Interfaces Configuration Statements and Hierarchy	75
	[edit firewall] Hierarchy Level	75
	[edit interfaces] Hierarchy Level	76
	[edit logical-systems] Hierarchy Level	92
	[edit protocols pppoe] Hierarchy Level	97
Chapter 4	Statement Summary	99
	address	102
	aggregate (Hierarchical Policer)	104
	accounting	105
	arp (Interfaces)	106
	bandwidth-limit (Hierarchical Policer)	107
	broadcast	108
	bundle	109
	burst-size-limit (Hierarchical Policer)	110
	cbr	111
	demux0 (Dynamic Interface)	112
	destination (IPCP)	113
	destination (Tunnels)	114
	destination-class-usage	115
	destination-profile	115
	dynamic-profiles	116
	epd-threshold (Logical Interface)	123
	eui-64	124
	family (Dynamic Standard Interface)	125
	family	127
	fast-aps-switch	131
	filter	132
	forward-and-send-to-re	133
	forward-only	133
	hierarchical-policer	134
	if-exceeding	135
	input	135
	input-list	136
	interface-mode	137
	interfaces	138
	interfaces (Static and Dynamic Subscribers)	139
	inverse-arp	143

ipsec-sa	143
keep-address-and-control	144
logical-systems	144
mode (Dynamic Profiles)	145
mode (Interfaces)	145
mtu	146
multicast-only	148
multipoint-destination	149
negotiate-address	150
no-redirects	150
oam-liveness	151
oam-period	152
output	153
output-list	153
policer (Interface)	154
post-service-filter	155
preferred	155
preferred-source-address	156
premium (Hierarchical Policer)	157
primary (Address on Interface)	158
protocols	158
proxy	159
queue-length	160
receive-options-packets	160
receive-ttl-exceeded	161
remote	161
rpf-check (Dynamic Profiles)	162
rpf-check (interfaces)	163
rtvbr	164
sampling (Interfaces)	165
service (Logical Interfaces)	166
service-filter (Interfaces)	167
service-set	167
shaping	168
source-class-usage	169
targeted-broadcast	170
then	171
translate-discard-eligible	171
translate-fecn-and-becn	172
unit (Dynamic Profiles Standard Interface)	173
unit	176
unnumbered-address (Demux)	182
unnumbered-address (Dynamic Profiles)	183
unnumbered-address (Ethernet)	185
unnumbered-address (PPP)	186
vbr	187
vci	188
vlan-id (Logical Port in Bridge Domain)	189
vlan-id-list (Interface in Bridge Domain)	189

Part 3	Administration	
Chapter 5	Monitoring Commands	193
	show interfaces (10-Gigabit Ethernet)	194
	show interfaces (Fast Ethernet)	219
	show interfaces (Gigabit Ethernet)	235
Chapter 6	ANCP Operational Mode Commands	259
	clear ancp neighbor	260
	clear ancp statistics	262
	clear ancp subscriber	264
	request ancp oam interface	266
	request ancp oam neighbor	267
	show ancp cos	268
	show ancp neighbor	271
	show ancp statistics	279
	show ancp subscriber	284
	show ancp summary	289
	show ancp summary neighbor	291
	show ancp summary subscriber	293
Chapter 7	BFD Operational Mode Commands	295
	clear bfd adaptation	296
	clear bfd session	297
	show bfd session	298
Chapter 8	BGP Operational Mode Commands	307
	clear bgp damping	308
	clear bgp neighbor	309
	clear bgp table	311
	show bgp bmp	313
	show bgp group	315
	show bgp group traffic-statistics	322
	show bgp neighbor	324
	show bgp replication	338
	show bgp summary	341
	show policy damping	347
Chapter 9	ES-IS Operational Mode Commands	349
	clear esis adjacency	350
	clear esis statistics	351
	show esis adjacency	352
	show esis interface	354
	show esis statistics	356
Chapter 10	IP Multicast Operational Mode Commands	359
	clear amt statistics	362
	clear amt tunnel	363
	clear igmp membership	364
	clear igmp snooping membership	367
	clear igmp snooping statistics	368

clear igmp statistics	369
clear mld membership	371
clear mld statistics	372
clear msdp cache	373
clear msdp statistics	374
clear multicast bandwidth-admission	375
clear multicast forwarding-cache	377
clear multicast scope	378
clear multicast sessions	379
clear multicast snooping statistics	380
clear multicast statistics	381
clear pgm negative-acknowledgments	382
clear pgm source-path-messages	383
clear pgm statistics	384
clear pim join	385
clear pim join-distribution	387
clear pim register	389
clear pim snooping join	391
clear pim snooping statistics	393
clear pim statistics	395
request pim multicast-tunnel rebalance	397
show amt statistics	398
show amt summary	401
show amt tunnel	403
show dvmrp interfaces	406
show dvmrp neighbors	408
show dvmrp prefix	410
show dvmrp prunes	412
show igmp group	414
show igmp interface	418
show igmp snooping interface	422
show igmp snooping membership	425
show igmp snooping statistics	429
show igmp statistics	432
show mld group	435
show mld interface	439
show mld statistics	443
show msdp	446
show msdp source	448
show msdp source-active	450
show msdp statistics	453
show multicast backup-pe-groups	457
show multicast flow-map	459
show multicast forwarding-cache statistics	461
show multicast interface	463
show multicast mrinfo	465
show multicast next-hops	467
show multicast pim-to-igmp-proxy	470
show multicast pim-to-mld-proxy	472

	show multicast route	474
	show multicast rpf	481
	show multicast scope	485
	show multicast snooping next-hops	487
	show multicast sessions	489
	show multicast snooping route	492
	show multicast statistics	495
	show multicast usage	498
	show pgm negative-acknowledgments	501
	show pgm source-path-messages	503
	show pgm statistics	504
	show pim bootstrap	507
	show pim interfaces	509
	show pim join	512
	show pim mdt	533
	show pim mdt data-mdt-joins	537
	show pim mdt data-mdt-limit	539
	show pim neighbors	541
	show pim rps	545
	show pim snooping interfaces	552
	show pim snooping join	555
	show pim snooping neighbors	559
	show pim snooping statistics	564
	show pim source	569
	show pim statistics	572
	show sap listen	585
	test msdp	587
Chapter 11	IPv6 Operational Mode Commands	589
	clear ipv6 neighbors	590
	clear ipv6 router-advertisement	591
	show ipv6 neighbors	592
	show ipv6 router-advertisement	594
Chapter 12	IS-IS Operational Mode Commands	597
	clear isis adjacency	598
	clear isis database	600
	clear isis overload	602
	clear isis statistics	604
	show isis adjacency	606
	show isis authentication	610
	show isis backup coverage	612
	show isis backup label-switched-path	614
	show isis backup spf results	616
	show isis context-identifier	620
	show isis database	622
	show isis hostname	634
	show isis interface	636
	show isis overview	640
	show isis route	643

	show isis spf	647
	show isis statistics	652
Chapter 13	LLDP Operational Mode Commands	655
	clear lldp neighbor	656
	clear lldp statistics	657
	show lldp	658
	show lldp local-information	661
	show lldp neighbors	664
	show lldp remote-global-statistics	669
	show lldp statistics	671
Chapter 14	MVRP Operational Mode Commands	673
	show mvrp	674
	show mvrp applicant-state	676
	show mvrp dynamic-vlan-memberships	678
	show mvrp interface	679
	show mvrp registration-state	680
	show mvrp statistics	682
Chapter 15	OSPF Operational Mode Commands	685
	clear (ospf ospf3) database	686
	clear (ospf ospf3) io-statistics	689
	clear (ospf ospf3) neighbor	690
	clear (ospf ospf3) overload	692
	clear (ospf ospf3) statistics	693
	show (ospf ospf3) backup coverage	695
	show (ospf ospf3) backup lsp	698
	show (ospf ospf3) backup neighbor	700
	show (ospf ospf3) backup spf	702
	show ospf context-identifier	709
	show ospf database	711
	show ospf3 database	719
	show (ospf ospf3) interface	730
	show (ospf ospf3) io-statistics	736
	show (ospf ospf3) log	738
	show (ospf ospf3) neighbor	741
	show (ospf ospf3) overview	747
	show (ospf ospf3) route	752
	show (ospf ospf3) statistics	758
Chapter 16	Protocol-Independent Routing Operational Mode Commands	763
	show as-path	765
	show as-path domain	769
	show as-path summary	771
	show route	773
	show route active-path	778
	show route advertising-protocol	783
	show route all	788
	show route aspath-regex	790

	show route best	792
	show route brief	795
	show route ccc	797
	show route community	798
	show route community-name	800
	show route damping	802
	show route detail	807
	show route exact	824
	show route export	826
	show route export vrf-target	829
	show route extensive	831
	show route flow validation	848
	show route forwarding-table	850
	show route forwarding-table interface-name	864
	show route hidden	867
	show route inactive-path	870
	show route inactive-prefix	873
	show route instance	875
	show route label	882
	show route label-switched-path	885
	show route localization	887
	show route martians	889
	show route match-prefix	891
	show route next-hop	892
	show route no-community	898
	show route output	901
	show route protocol	906
	show route range	918
	show route receive-protocol	922
	show route resolution	930
	show route snooping	933
	show route source-gateway	941
	show route summary	947
	show route table	951
	show route terse	965
Chapter 17	RIP Operational Mode Commands	969
	clear rip general-statistics	970
	clear rip statistics	971
	show rip general-statistics	972
	show rip neighbor	974
	show rip statistics	976
Chapter 18	RIPng Operational Mode Commands	979
	clear ripng general-statistics	980
	clear ripng statistics	981
	show ripng general-statistics	982
	show ripng neighbor	983
	show ripng statistics	985

Chapter 19	Firewall Filter Operational Mode Commands	987
	clear firewall	988
	show firewall	990
	show firewall filter version	997
	show firewall log	998
	show firewall prefix-action-stats	1001
	show firewall templates-in-use	1003
	show policer	1005
Chapter 20	Layer 2 Bridging and Switching Operational Mode Commands	1009
	clear bridge mac-table	1010
	clear error bpdu	1011
	clear error mac-rewrite	1012
	show bridge domain	1013
	show bridge flood	1015
	show bridge mac-table	1022
	show bridge statistics	1026
	show l2-learning global-information	1029
	show l2-learning global-mac-count	1031
	show l2-learning instance	1032
	show l2-learning interface	1034
	show mac-rewrite interface	1036
Chapter 21	VPN Operational Mode Commands	1037
	clear vpls mac-address	1038
	clear vpls mac-table	1039
	request l2circuit-switchover	1040
	show dynamic-tunnels database	1041
	show hfrf profiles	1044
	show ingress-replication mvpn	1046
	show l2circuit connections	1048
	show l2vpn connections	1055
	show mvpn c-multicast	1061
	show mvpn instance	1063
	show mvpn neighbor	1067
	show vpls connections	1071
	show vpls flood event-queue	1082
	show vpls flood instance	1084
	show vpls flood route	1086
	show vpls mac-table	1088
	show vpls statistics	1093
Part 4	Troubleshooting	
Chapter 22	Interface Diagnostics	1099
	Interface Diagnostics	1099
	Configuring Loopback Testing	1099
	Interface Diagnostics	1101
	Starting and Stopping a BERT Test	1105
	Example: Configuring Bit Error Rate Testing	1105

Part 5**Index**

Index	1109
-------------	------

List of Figures

Part 2	Configuration	
Chapter 2	Protocol Family and Interface Address Properties	9
	Figure 1: Hierarchical Policer	15
	Figure 2: Prefix Accounting with Source and Destination Classes	26
	Figure 3: Unicast RPF with Routing Asymmetry	54
	Figure 4: Unicast RPF Sample Topoolgy	66

List of Tables

	About the Documentation	xxiii
	Table 1: Notice Icons	xxvi
	Table 2: Text and Syntax Conventions	xxvi
Part 3	Administration	
Chapter 5	Monitoring Commands	193
	Table 3: show interfaces Gigabit Ethernet Output Fields	195
	Table 4: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type . .	209
	Table 5: show interfaces Fast Ethernet Output Fields	219
	Table 6: show interfaces Gigabit Ethernet Output Fields	236
	Table 7: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type . .	249
Chapter 6	ANCP Operational Mode Commands	259
	Table 8: show ancp cos Output Fields	268
	Table 9: show ancp neighbor Output Fields	271
	Table 10: show ancp statistics Output Fields	279
	Table 11: show ancp subscriber Output Fields	284
	Table 12: show ancp summary Output Fields	289
	Table 13: show ancp summary neighbor Output Fields	291
	Table 14: show ancp summary subscriber Output Fields	293
Chapter 7	BFD Operational Mode Commands	295
	Table 15: show bfd session Output Fields	299
Chapter 8	BGP Operational Mode Commands	307
	Table 16: show bgp bmp Output Fields	313
	Table 17: show bgp group Output Fields	316
	Table 18: show bgp group traffic-statistics Output Fields	322
	Table 19: show bgp neighbor Output Fields	325
	Table 20: show bgp replication Output Fields	338
	Table 21: show bgp summary Output Fields	342
	Table 22: show policy damping Output Fields	347
Chapter 9	ES-IS Operational Mode Commands	349
	Table 23: show esis adjacency Output Fields	352
	Table 24: show esis interface Output Fields	354
	Table 25: show esis statistics Output Fields	356
Chapter 10	IP Multicast Operational Mode Commands	359
	Table 26: show amt statistics Output Fields	398
	Table 27: show amt summary Output Fields	401
	Table 28: show amt tunnel Output Fields	403

Table 29: show dvmrp interfaces Output Fields	406
Table 30: show dvmrp neighbors Output Fields	408
Table 31: show dvmrp prefix Output Fields	410
Table 32: show dvmrp prunes Output Fields	412
Table 33: show igmp group Output Fields	414
Table 34: show igmp interface Output Fields	418
Table 35: show igmp snooping interface Output Fields	422
Table 36: show igmp snooping membership Output Fields	425
Table 37: show igmp snooping statistics Output Fields	429
Table 38: show igmp statistics Output Fields	432
Table 39: show mld group Output Fields	435
Table 40: show mld interface Output Fields	439
Table 41: show mld statistics Output Fields	443
Table 42: show msdp Output Fields	446
Table 43: show msdp source Output Fields	449
Table 44: show msdp source-active Output Fields	451
Table 45: show msdp statistics Output Fields	453
Table 46: show multicast backup-pe-groups Output Fields	457
Table 47: show multicast flow-map Output Fields	459
Table 48: show multicast forwarding-cache statistics Output Fields	461
Table 49: show multicast interface Output Fields	463
Table 50: show multicast minfo Output Fields	465
Table 51: show multicast next-hops Output Fields	468
Table 52: show multicast pim-to-igmp-proxy Output Fields	470
Table 53: show multicast pim-to-mld-proxy Output Fields	472
Table 54: show multicast route Output Fields	475
Table 55: show multicast rpf Output Fields	482
Table 56: show multicast scope Output Fields	485
Table 57: show multicast snooping next-hops Output Fields	487
Table 58: show multicast sessions Output Fields	489
Table 59: show multicast snooping route Output Fields	493
Table 60: show multicast statistics Output Fields	495
Table 61: show multicast usage Output Fields	498
Table 62: show pgm negative-acknowledgments Output Fields	501
Table 63: show pgm source-path-messages Output Fields	503
Table 64: show pgm statistics Output Fields	504
Table 65: show pim bootstrap Output Fields	507
Table 66: show pim interfaces Output Fields	509
Table 67: show pim join Output Fields	514
Table 68: show pim mdt Output Fields	533
Table 69: show pim mdt data-mdt-joins Output Fields	537
Table 70: show pim mdt data-mdt-limit Output Fields	539
Table 71: show pim neighbors Output Fields	542
Table 72: show pim rps Output Fields	546
Table 73: show pim snooping interface Output Fields	552
Table 74: show pim snooping join Output Fields	555
Table 75: show pim snooping neighbors Output Fields	559
Table 76: show pim snooping statistics Output Fields	564
Table 77: show pim source Output Fields	570

	Table 78: show pim statistics Output Fields	573
	Table 79: show sap listen Output Fields	585
Chapter 11	IPv6 Operational Mode Commands	589
	Table 80: show ipv6 neighbors Output Fields	592
	Table 81: show ipv6 router-advertisement Output Fields	594
Chapter 12	IS-IS Operational Mode Commands	597
	Table 82: show isis adjacency Output Fields	606
	Table 83: show isis authentication Output Fields	610
	Table 84: show isis backup coverage Output Fields	612
	Table 85: show isis backup label-switched-path Output Fields	614
	Table 86: show isis backup spf results Output Fields	617
	Table 87: show isis context-identifier Output Fields	620
	Table 88: show isis database Output Fields	623
	Table 89: show isis hostname Output Fields	634
	Table 90: show isis interface Output Fields	637
	Table 91: show isis overview Output Fields	640
	Table 92: show isis route Output Fields	643
	Table 93: show isis spf Output Fields	647
	Table 94: show isis statistics Output Fields	653
Chapter 13	LLDP Operational Mode Commands	655
	Table 95: show lldp Output Fields	658
	Table 96: show lldp local-information Output Fields	661
	Table 97: show lldp neighbors Output Fields	664
	Table 98: show lldp remote-global-statistics Output Fields	669
	Table 99: show lldp statistics Output Fields	671
Chapter 14	MVRP Operational Mode Commands	673
	Table 100: show mvrp Output Fields	674
	Table 101: show mvrp applicant-state Output Fields	676
	Table 102: show mvrp dynamic-vlan-memberships Output Fields	678
	Table 103: show mvrp interface Output Fields	679
	Table 104: show mvrp registration-state Output Fields	680
	Table 105: show mvrp statistics Output Fields	682
Chapter 15	OSPF Operational Mode Commands	685
	Table 106: show (ospf ospf3) backup coverage Output Fields	695
	Table 107: show (ospf ospf3) backup lsp Output Fields	698
	Table 108: show (ospf ospf3) backup neighbor Output Fields	700
	Table 109: show (ospf ospf3) backup spf Output Fields	703
	Table 110: show ospf context-identifier Output Fields	710
	Table 111: show ospf database Output Fields	712
	Table 112: show ospf3 database Output Fields	720
	Table 113: show (ospf ospf3) interface Output Fields	731
	Table 114: show (ospf ospf3) io-statistics Output Fields	736
	Table 115: show (ospf ospf3) log Output Fields	738
	Table 116: show (ospf ospf3) neighbor Output Fields	742
	Table 117: show ospf overview Output Fields	747
	Table 118: show (ospf ospf3) route Output Fields	753

	Table 119: show (ospf ospf3) statistics Output Fields	758
Chapter 16	Protocol-Independent Routing Operational Mode Commands	763
	Table 120: show as-path Output Fields	765
	Table 121: show as-path domain Output Fields	769
	Table 122: show as-path summary Output Fields	771
	Table 123: show route Output Fields	774
	Table 124: show route advertising-protocol Output Fields	784
	Table 125: show route damping Output Fields	802
	Table 126: show route detail Output Fields	807
	Table 127: Next-hop Types Output Field Values	812
	Table 128: State Output Field Values	813
	Table 129: Communities Output Field Values	815
	Table 130: show route export Output Fields	826
	Table 131: show route export vrf-target Output Fields	829
	Table 132: show route extensive Output Fields	831
	Table 133: show route flow validation Output Fields	848
	Table 134: show route forwarding-table Output Fields	853
	Table 135: show route forwarding-table interface-name Output Fields	864
	Table 136: show route instance Output Fields	876
	Table 137: show route localization Output Fields	887
	Table 138: show route martians Output Fields	889
	Table 139: show route receive-protocol Output Fields	922
	Table 140: show route resolution Output Fields	931
	Table 141: show route summary Output Fields	947
	Table 142: show route terse Output Fields	965
Chapter 17	RIP Operational Mode Commands	969
	Table 143: show rip general-statistics Output Fields	972
	Table 144: show rip neighbor Output Fields	974
	Table 145: show rip statistics Output Fields	977
Chapter 18	RIPng Operational Mode Commands	979
	Table 146: show ripng general-statistics Output Fields	982
	Table 147: show ripng neighbor Output Fields	983
	Table 148: show ripng statistics Output Fields	985
Chapter 19	Firewall Filter Operational Mode Commands	987
	Table 149: show firewall Output Fields	991
	Table 150: show firewall filter version Output Fields	997
	Table 151: show firewall log Output Fields	998
	Table 152: show firewall prefix-action-stats Output Fields	1001
	Table 153: show firewall templates-in-use Output Fields	1003
	Table 154: show policer Output Fields	1005
Chapter 20	Layer 2 Bridging and Switching Operational Mode Commands	1009
	Table 155: show bridge mac-table Output fields	1023
	Table 156: show l2-learning global-information Output Fields	1029
	Table 157: show l2-learning instance Output Fields	1032
	Table 158: show l2-learning interfaceOutput Fields	1034
	Table 159: show mac-rewrite interface Output Fields	1036

Chapter 21	VPN Operational Mode Commands	1037
	Table 160: show dynamic-tunnels database Output Fields	1041
	Table 161: show hfrp profiles Output Fields	1044
	Table 162: show ingress-replication mvpn	1046
	Table 163: show l2circuit connections Output Fields	1049
	Table 164: show l2vpn connections Output Fields	1056
	Table 165: show mvpn c-multicast Output Fields	1061
	Table 166: show mvpn instance Output Fields	1063
	Table 167: show mvpn neighbor Output Fields	1067
	Table 168: show vpls connections Output Fields	1072
	Table 169: show vpls flood event-queue Output Fields	1082
	Table 170: show vpls flood instance Output Fields	1084
	Table 171: show vpls flood route Output Fields	1086
	Table 172: show vpls mac-table Output fields	1088
	Table 173: show vpls statistics Output Fields	1093
 Part 4	 Troubleshooting	
Chapter 22	Interface Diagnostics	1099
	Table 174: Loopback Modes by Interface Type	1100
	Table 175: BERT Capabilities by Interface Type	1104

About the Documentation

- Documentation and Release Notes on page xxiii
- Supported Platforms on page xxiii
- Using the Examples in This Manual on page xxiv
- Documentation Conventions on page xxv
- Documentation Feedback on page xxvii
- Requesting Technical Support on page xxviii

Documentation and Release Notes

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

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

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

Supported Platforms

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

- ACX Series
- M Series
- MX Series
- T Series
- J Series
- PTX Series

Using the Examples in This Manual

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

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

Documentation Conventions

[Table 1 on page xxvi](#) defines notice icons used in this guide.

Table 1: Notice Icons

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

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

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies guide names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS CLI User Guide</i> RFC 1997, <i>BGP Communities Attribute</i>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>

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

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

Documentation Feedback

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

- Document or topic name
- URL or page number

- Software release version (if applicable)

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

- Find CSC offerings: <http://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>
- Find product documentation: <http://www.juniper.net/techpubs/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>
- Download the latest versions of software and review release notes: <http://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <http://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <http://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <http://www.juniper.net/cm/>

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

Opening a Case with JTAC

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

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

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

PART 1

Overview

- [Protocol Family and Interface Address Properties on page 3](#)

CHAPTER 1

Protocol Family and Interface Address Properties

- [Protocol Family Configuration and Interface Address Statements on page 3](#)

Protocol Family Configuration and Interface Address Statements

For each logical interface, you must configure one or more protocol families. You can also configure interface address properties. To do this, include the following statements:

```
family family {  
    accounting {  
        destination-class-usage;  
        source-class-usage {  
            direction;  
        }  
    }  
    address address {  
        destination address;  
    }  
    bundle interface-name;  
    filter {  
        dialer filter-name;  
        input filter-name;  
        output filter-name;  
        group filter-group-number;  
    }  
    interface-mode (access | trunk);  
    ipsec-sa sa-name;  
    keep-address-and-control;  
    mtu bytes;  
    multicast-only;  
    negotiate-address;  
    no-redirects;  
    policer {  
        arp policer-template-name;  
        input policer-template-name;  
        output policer-template-name;  
    }  
    primary;  
    protocols [inet iso mpls];  
    proxy inet-address address;
```

```
receive-options-packets;
receive-ttl-exceeded;
remote (inet-address address | mac-address address);
rpf-check <fail-filter filter-name>;
sampling {
    direction;
}
service {
    input {
        service-set service-set-name <service-filter filter-name>;
        post-service-filter filter-name;
    }
    output {
        service-set service-set-name <service-filter filter-name>;
    }
}
targeted-broadcast {
    forward-and-send-to-re;
    forward-only;
}
(translate-discard-eligible | no-translate-discard-eligible);
(translate-fecn-and-becn | no-translate-fecn-and-becn);
translate-plp-control-word-de;
vlan-id number;
vlan-id-list [number number-number];
unnumbered-address interface-name destination address destination-profile
    profile-name;
address address {
    arp ip-address (mac | multicast-mac) mac-address <publish>;
    broadcast address;
    destination address;
    destination-profile name;
    eui-64;
    multipoint-destination address dlcid dlcid-identifier;
    multipoint-destination address {
        epd-threshold cells;
        inverse-arp;
        oam-liveness {
            up-count cells;
            down-count cells;
        }
        oam-period (disable | seconds);
        shaping {
            (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate
                burst length);
            queue-length number;
        }
        vci vpi-identifier.vci-identifier;
    }
    primary;
    preferred;
    (vrrp-group | vrrp-inet6-group) group-number {
        (accept-data | no-accept-data);
        advertise-interval seconds;
        authentication-type authentication;
        authentication-key key;
    }
}
```

```

    fast-interval milliseconds;
    (preempt | no-preempt) {
        hold-time seconds;
    }
    priority-number number;
    track {
        priority-cost seconds;
        priority-hold-time interface-name {
            interface priority;
            bandwidth-threshold bits-per-second {
                priority;
            }
        }
        route ip-address/mask routing-instance instance-name priority-cost cost;
    }
    virtual-address [ addresses ];
}
}

```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* **unit** *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* **unit** *logical-unit-number*]

For information about interface-specific protocol and address properties, see *Configuring T1 and NxDSO Interfaces*.

Related Documentation

- *Configuring T1 and NxDSO Interfaces*

PART 2

Configuration

- [Protocol Family and Interface Address Properties on page 9](#)
- [Network Interfaces Configuration Statements and Hierarchy on page 75](#)
- [Statement Summary on page 99](#)

CHAPTER 2

Protocol Family and Interface Address Properties

- [Setting the Protocol MTU on page 9](#)
- [Disabling the Removal of Address and Control Bytes on page 10](#)
- [Disabling the Transmission of Redirect Messages on an Interface on page 10](#)
- [Applying Policers on page 11](#)
- [Applying a Filter to an Interface on page 20](#)
- [Enabling Source Class and Destination Class Usage on page 25](#)
- [Configuring the Protocol Family on page 33](#)
- [Configuring the Interface Address on page 36](#)
- [Configuring ICCP for MC-LAG on page 40](#)
- [Configuring IPCP Options on page 42](#)
- [Configuring an Unnumbered Interface on page 44](#)
- [Configuring Default, Primary, and Preferred Addresses and Interfaces on page 50](#)
- [Configuring Unicast RPF on page 52](#)
- [Configuring Targeted Broadcast on page 61](#)
- [Understanding Targeted Broadcast on page 63](#)
- [Example: Configuring Unicast Reverse-Path-Forwarding Check on page 64](#)

Setting the Protocol MTU

When you initially configure an interface, the protocol maximum transmission unit (MTU) is calculated automatically. If you subsequently change the media MTU, the protocol MTU on existing address families automatically changes.

For a list of default protocol MTU values, see *Configuring the Media MTU*.

To modify the MTU for a particular protocol family, include the **mtu** statement:

mtu *bytes*;

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number* family *family*]**

- `[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family family]`

If you increase the size of the protocol MTU, you must ensure that the size of the media MTU is equal to or greater than the sum of the protocol MTU and the encapsulation overhead. For a list of encapsulation overhead values, see *Configuring the Media MTU*. If you reduce the media MTU size, but there are already one or more address families configured and active on the interface, you must also reduce the protocol MTU size. (You configure the media MTU by including the `mtu` statement at the `[edit interfaces interface-name]` hierarchy level, as discussed in *Configuring the Media MTU*.)



NOTE: Changing the media MTU or protocol MTU causes an interface to be deleted and added again.

The maximum number of data-link connection identifiers (DLCIs) is determined by the MTU on the interface. If you have keepalives enabled, the maximum number of DLCIs is 1000, with the MTU set to 5012.

The actual frames transmitted also contain cyclic redundancy check (CRC) bits, which are not part of the MTU. For example, the default protocol MTU for a Gigabit Ethernet interface is 1500 bytes, but the largest possible frame size is actually 1504 bytes; you need to consider the extra bits in calculations of MTUs for interoperability.

**Related
Documentation**

- *Configuring the Media MTU*

Disabling the Removal of Address and Control Bytes

For Point-to-Point Protocol (PPP) CCC-encapsulated interfaces, the address and control bytes are removed by default before the packet is encapsulated into a tunnel.

You can disable the removal of address and control bytes. To do this, include the `keep-address-and-control` statement:

```
keep-address-and-control;
```

You can include this statement at the following hierarchy levels:

- `[edit interfaces interface-name unit logical-unit-number family ccc]`
- `[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family ccc]`

**Related
Documentation**

- `keep-address-and-control` on page 144

Disabling the Transmission of Redirect Messages on an Interface

By default, the interface sends protocol redirect messages. To disable the sending of these messages on an interface, include the `no-redirects` statement:

no-redirects;

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number* family *family*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*]**

To disable the sending of protocol redirect messages for the entire router or switch, include the **no-redirects** statement at the **[edit system]** hierarchy level.

Related
Documentation

- [no-redirects on page 150](#)

Applying Policers

Policies allow you to perform simple traffic policing on specific interfaces or Layer 2 virtual private networks (VPNs) without configuring a firewall filter. To apply policies, include the **policer** statement:

```
policer {
  arp policer-template-name;
  disable-arp-policer
  input policer-template-name;
  output policer-template-name;
}
```

You can include these statements at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number* family *family*]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*]**

In the **family** statement, the protocol family can be **ccc**, **inet**, **inet6**, **mpls**, **tcc**, or **vpls**.

In the **arp** statement, list the name of one policer template to be evaluated when Address Resolution Protocol (ARP) packets are received on the interface. By default, an ARP policer is installed that is shared among all the Ethernet interfaces on which you have configured the **family inet** statement. If you want more stringent or lenient policing of ARP packets, you can configure an interface-specific policer and apply it to the interface. You configure an ARP policer just as you would configure any other policer, at the **[edit firewall policer]** hierarchy level. If you apply this policer to an interface, the default ARP packet policer is overridden. If you delete this policer, the default policer takes effect again.

Starting with Junos OS Release 14.1, you can configure the router to disable the processing of the specified ARP policies on the received ARP packets. Disabling ARP policies can cause denial-of-service (DoS) attacks on the system. Due to this possibility, we recommend that you exercise caution while disabling ARP policies. To prevent the processing of ARP policies on the arriving ARP packets, include the **disable-arp-policer** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family *inet* policer]** or the **[edit logical-systems *logical-system-name* interfaces *interface-name* unit**

logical-unit-number family inet policer] hierarchy level. You can configure this statement only for interfaces with inet address families and on MX Series routers with MPCs. When you disable ARP policers per interface, the packets are continued to be policed by the distributed DoS (DDoS) ARP policer. The maximum rate of is 10000 pps per FPC.

In the **input** statement, list the name of one policer template to be evaluated when packets are received on the interface.

In the **output** statement, list the name of one policer template to be evaluated when packets are transmitted on the interface.



NOTE: To use policing on a CCC or TCC interface, you must configure the CCC or TCC protocol family.

You can configure a different policer on each protocol family on an interface, with one input policer and one output policer for each family. When you apply policers, you can configure the family **ccc**, **inet**, **inet6**, **mpls**, **tcc**, or **vpls** only, and one ARP policer for the family **inet** protocol only. Each time a policer is referenced, a separate copy of the policer is installed on the packet forwarding components for that interface.

If you apply both policers and firewall filters to an interface, input policers are evaluated before input firewall filters, and output policers are evaluated after output firewall filters.

If you apply the policer to the interface **lo0**, it is applied to packets received or transmitted by the Routing Engine.

On T Series, M120, and M320 platforms, if the interfaces are on the same FPC, the filters or policers do not act on the sum of traffic entering and exiting the interfaces.

For more information about policers, see the *Routing Policy Feature Guide for Routing Devices*.

This section includes the following topics:

- [Applying Aggregate Policers on page 12](#)
- [Applying Hierarchical Policers on Enhanced Intelligent Queuing PICs on page 14](#)

Applying Aggregate Policers

By default, if you apply a policer to multiple protocol families on the same logical interface, the policer restricts traffic for each protocol family individually. For example, a policer with a 50 Mbps bandwidth limit applied to both IPv4 and IPv6 traffic would allow the interface to accept 50 Mbps of IPv4 traffic and 50 Mbps of IPv6 traffic. If you apply an aggregate policer, the policer would allow the interface to receive only 50 Mbps of IPv4 and IPv6 traffic combined.

To configure an aggregate policer, include the **logical-interface-policer** statement at the **[edit firewall policer policer-template-name]** hierarchy level:

```
[edit firewall policer policer-template-name]  
logical-interface-policer;
```

For the policer to be treated as an aggregate, you must apply it to multiple protocol families on a single logical interface by including the **policer** statement:

```
policer {
  arp policer-template-name;
  disable-arp-policer
  input policer-template-name;
  output policer-template-name;
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family *family*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*]

In the **family** statement, the protocol family can be **ccc**, **inet**, **inet6**, **mpls**, **tcc**, or **vpls**.

The protocol families on which you do not apply the policer are not affected by the policer. For example, if you configure a single logical interface to accept MPLS, IPv4, and IPv6 traffic and you apply the logical interface policer **policer1** to only the IPv4 and IPv6 protocol families, MPLS traffic is not subject to the constraints of **policer1**.

If you apply **policer1** to a different logical interface, there are two instances of the policer. This means the Junos OS polices traffic on separate logical interfaces separately, not as an aggregate, even if the same logical-interface policer is applied to multiple logical interfaces on the same physical interface port.

Example: Applying Aggregate Policers

Configure two logical interface policers: **aggregate_police1** and **aggregate_police2**. Apply **aggregate_police1** to IPv4 and IPv6 traffic received on logical interface **fe-0/0/0.0**. Apply **aggregate_police2** to CCC and MPLS traffic received on logical interface **fe-0/0/0.0**. This configuration causes the software to create only one instance of **aggregate_police1** and one instance of **aggregate_police2**.

Apply **aggregate_police1** to IPv4 and IPv6 traffic received on another logical interface **fe-0/0/0.1**. This configuration causes the software to create a new instance of **aggregate_police1**, one that applies to unit 0 and another that applies to unit 1.

```
[edit firewall]
policer aggregate_police1 {
  logical-interface-policer;
  if-exceeding {
    bandwidth-limit 100m;
    burst-size-limit 500k;
  }
  then {
    discard;
  }
}
policer aggregate_police2 {
  logical-interface-policer;
  if-exceeding {
```

```
        bandwidth-limit 10m;
        burst-size-limit 200k;
    }
    then {
        discard;
    }
}
[edit interfaces fe-0/0/0]
unit 0 {
    family inet {
        policer {
            input aggregate_police1;
        }
    }
    family inet6 {
        policer {
            input aggregate_police1;
        }
    }
    family ccc {
        policer {
            input aggregate_police2;
        }
    }
    family mpls {
        policer {
            input aggregate_police2;
        }
    }
}
unit 1 {
    family inet {
        policer {
            input aggregate_police1;
        }
    }
    family inet6 {
        policer {
            input aggregate_police1;
        }
    }
}
```

Applying Hierarchical Policers on Enhanced Intelligent Queuing PICs

M40e, M120, and M320 edge routers and T Series core routers with Enhanced Intelligent Queuing (IQE) PICs support hierarchical policers in the ingress direction and allow you to apply a hierarchical policer for the premium and aggregate (premium plus normal) traffic levels to an interface. Hierarchical policers provide cross-functionality between the configured physical interface and the Packet Forwarding Engine.


Before you begin, there are some general restrictions that apply to hierarchical policers:

- Only one type of policer can be configured for a logical or physical interface. For example, a hierarchical policer and a regular policer in the same direction for the same logical interface is not allowed.
- The chaining of the policers—that is, applying policers to both a port and the logical interfaces of that port—is not allowed.
- There is a limit of 64 policers per interface in case there is no BA classification, providing a single policer per DLCI.
- Only one kind of policer can be applied on a physical or logical interface.
- The policer should be independent of BA classification. Without BA classification, all traffic on an interface will be treated either as EF or non-EF, based on the configuration. With BA classification, an interface can support up to 64 policers. Again, the interface here may be a physical interface or logical interface (for example, DLCI).
- With BA classification, the miscellaneous traffic (the traffic *not* matching with any of the BA classification DSCP/EXP bits) will be policed as non-EF traffic. No separate policers will be installed for this traffic.

Hierarchical Policer Overview

Hierarchical policing uses two token buckets, one for aggregate (non-EF) traffic and one for premium (EF) traffic. Which traffic is EF and which is non-EF is determined by the class-of-service configuration. Logically, hierarchical policing is achieved by chaining two policers.

Figure 1: Hierarchical Policer



In the example in [Figure 1 on page 15](#), EF traffic is policed by Premium Policer and non EF traffic is policed by Aggregate Policer. What that means is, for EF traffic the out-of-spec action will be the one that is configured for Premium Policer, but the in-spec EF traffic will still consume the tokens from the Aggregate Policer.

But EF traffic will never be submitted to the out-of-spec action of the Aggregate Policer. Also, if the out-of-spec action of the Premium Policer is not set to Discard, those out-of-spec packets will not consume the tokens from the Aggregate Policer. Aggregate Policer only polices the non-EF traffic. As you can see, the Aggregate Policer token bucket can go negative, if all the tokens are consumed by the non-EF traffic and then you get bursts of EF traffic. But that will be for a very short time, and over a period of time it will average out. For example:

- *Premium Policer*: Bandwidth 2 Mbps, OOS Action: Discard
- *Aggregate Policer*: Bandwidth 10 Mbps, OOS Action: Discard

In the above case, EF traffic is guaranteed 2 Mbps and the non-EF traffic will get from 8 Mbps to 10 Mbps, depending on the input rate of the EF traffic.

Hierarchical Policing Characteristics

Hierarchical token bucket features include:

- Ingress traffic is first classified into EF and non-EF traffic prior to applying a policer:
 - Classification is performed by Q-tree lookup
- Channel number selects a shared token bucket policer:
 - Dual token bucket policer is divided into two single bucket policers:
 - Policer1—EF traffic
 - Policer2—non-EF traffic
- Shared token bucket is used to police the traffic as follows:
 - Policer1 is set to EF rate (for example, 2 Mbps)
 - Policer2 is set to aggregate interface policed rate (for example, 10 Mbps).
 - EF traffic gets applied to Policer1.
 - If traffic is in-spec it is allowed to pass and decrement from both Policer1 and Policer2.
 - If traffic is out-of-spec it can be discarded or marked with a new FC or loss priority. Policer2 will not do anything with out-of-spec EF traffic.
 - Non-EF traffic gets applied only to Policer2.
 - If traffic is in-spec it is allowed to pass through and decremented Policer2.
 - If traffic is out-of-spec it is discarded or marked with a new FC or set with a new drop priority.
- Rate-limit the port speed to a desired rate at Layer 2
- Rate-limit the EF traffic
- Rate-limit the non-EF traffic
- Policing drops counted per color

Configuring Hierarchical Policers

To configure a hierarchical policer, apply the **policing-priority** statement to the proper forwarding class and configure a hierarchical policer for the aggregate and premium level. For more information about class of service, see the *Junos OS Class of Service Library for Routing Devices*.



NOTE: Hierarchical policers can only be configured on SONET physical interfaces hosted on an IQE PIC. Only aggregate and premium levels are supported.

CoS Configuration of Forwarding Classes for Hierarchical Policers

```
[edit class-of-service forwarding-classes]
class fc1 queue-num 0 priority high policing-priority premium;
class fc2 queue-num 1 priority low policing-priority normal;
class fc3 queue-num 2 priority low policing-priority normal;
class fc4 queue-num 3 priority low policing-priority normal;
```

For detailed information on class-of-service configuration and statements, see the *Junos OS Class of Service Library for Routing Devices*.

Firewall Configuration for Hierarchical Policers

```
[edit firewall hierarchical-policer foo]
aggregate {
  if-exceeding {
    bandwidth-limit 70m;
    burst-size-limit 1500;
  }
  then {
    discard;
  }
}
premium {
  if-exceeding {
    bandwidth-limit 50m;
    burst-size-limit 1500;
  }
  then {
    discard;
  }
}
```

You can apply the hierarchical policer as follows:

```
[edit interfaces so-0/1/0 unit 0 layer-2-policer]
input-hierarchical-policer foo;
```

You also have the option to apply the policer at the physical port level as follows:

```
[edit interfaces so-0/1/0 layer-2-policer]
input-hierarchical-policer foo;
```

Configuring a Single-Rate Two-Color Policer

You can configure a single-rate two-color policer as follows:

```
[edit firewall policer foo]
if-exceeding {
  bandwidth-limit 50m;
  burst-size-limit 1500;
}
then {
  discard;
}
```

You can apply the policer as follows:

```
[edit interfaces so-0/1/0 unit 0 layer-2-policer]
input-policer foo;
```

You also have the option to apply the policer at the physical port level as follows:

```
[edit interfaces so-0/1/0 layer-2-policer]
input-policer foo;
```

Configuring a Single-Rate Tricolor Policer

This section describes single-rate color blind and color aware policers.

Configuring a Single-Rate Color-Blind Policer

You can configure a single-rate color blind policer as follows:

```
[edit firewall three-color-policer foo]
single-rate {
  color-blind;
  committed-information-rate 50m;
  committed-burst-size 1500;
  excess-burst-size 1500;
}
```

You can apply the single-rate color blind policer as follows:

```
[edit interfaces so-0/1/0 unit 0 layer-2-policer]
input-three-color foo;
```

You also have the option to apply the policer at the physical port level as follows:

```
[edit interfaces so-0/1/0 layer-2-policer]
input-three-color foo;
```

Configuring a Single-Rate Color-Aware Policer

You can configure a single-rate color-aware policer as follows:

```
[edit firewall three-color-policer bar]
single-rate {
  color-aware;
  committed-information-rate 50m;
  committed-burst-size 1500;
  excess-burst-size 1500;
}
```

You can apply the single-rate color-aware policer as follows:

```
[edit interfaces so-0/1/0 unit 0 layer-2-policer]
input-three-color foo;
```

You also have the option to apply the policer at the physical port level as follows:

```
[edit interfaces so-0/1/0 layer-2-policer]
input-three-color bar;
```

Configuring a Two-Rate Tricolor Marker Policer

Ingress policing is implemented using a two-rate tricolor marker (trTCM). This is done with a dual token bucket (DTB) that maintains two rates, committed, and a peak. Egress static policing also uses a token bucket.

The token buckets perform the following ingress policing functions:

- (1K) trTCM - Dual token bucket (red, yellow, and green marking)
- Policing is based on Layer 2 packet size:
 - After +/- byte adjust offset
- Marking is color aware and color blind:
 - Color aware needs to have the color set by q-tree lookup based on:
 - ToS
 - EXP
- Programmable marking actions:
 - Color (red, yellow, green)
 - Drop based on color and congestion profile
- Policer is selected based on the arriving channel number:
 - Channel number LUT produces policer index and queue index
 - Multiple channels can share the same policer (LUT produces same policer index)
- Support ingress policing and trTCM at the following levels:
 - Queue
 - Logical interface (ifl/DLCI)
 - Physical interface (ifd)
 - Physical port (controller ifd)
 - Any combinations of logical interface, physical interface, and port
- Support percentage of interface speed and bits per second

Rate limits may be applied to selected queues on ingress and on predefined queues at egress. The token bucket operates in color aware and color blind modes (specified by RFC 2698).

Configuring a Color-Blind trTCM

```
[edit firewall three-color-policer foo]
two-rate {
  color-blind;
  committed-information-rate 50m;
  committed-burst-size 1500;
  peak-information-rate 100m;
  peak-burst-size 3k;
```

```
}
```

You can apply the three-color two-rate color-blind policer as follows:

```
[edit interfaces so-0/1/0 unit 0 layer-2-policer]
input-three-color foo;
```

You also have the option to apply the policer at the physical port level as follows:

```
[edit interfaces so-0/1/0 layer-2-policer]
input-three-color foo;
```

Configuring a Color-Aware trTCM

```
[edit firewall three-color-policer bar]
two-rate {
  color-aware;
  committed-information-rate 50m;
  committed-burst-size 1500;
  peak-information-rate 100m;
  peak-burst-size 3k;
}
```

You can apply the three-color two-rate color-aware policer as follows:

```
[edit interfaces so-0/1/0 unit 0 layer-2-policer]
input-three-color bar;
```

You also have the option to apply the policer at the physical port level as follows:

```
[edit interfaces so-0/1/0 layer-2-policer]
input-three-color bar;
```

Related Documentation

- *Junos OS Class of Service Library for Routing Devices*

Applying a Filter to an Interface

To apply firewall filters to an interface, include the **filter** statement:

```
filter {
  group filter-group-number;
  input filter-name;
  input-list [ filter-names ];
  output filter-name;
  output-list [ filter-names ];
}
```

To apply a single filter, include the **input** statement:

```
filter {
  input filter-name;
}
```

To apply a list of filters to evaluate packets received on an interface, include the **input-list** statement.

```
filter {
  input-list [ filter-names ];
}
```

Up to 16 filter names can be included in an input list.

To apply a list of filters to evaluate packets transmitted on an interface, include the **output-list** statement.

```
filter {
  output-list [ filter-names ];
}
```

When you apply filters using the **input-list** statement or the **output-list** statement, a new filter is created with the name *<interface-name>.<unit-direction>*. This filter is exclusively interface-specific.

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family *family*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*]

In the **family** statement, the protocol family can be **ccc**, **inet**, **inet6**, **mpls**, or **vpls**.

In the **group** statement, specify the interface group number to associate with the filter.

In the **input** statement, list the name of one firewall filter to be evaluated when packets are received on the interface.

In the **input-list** statement, list the names of filters to evaluate when packets are received on the interface. You can include up to 16 filter names.

In the **output** statement, list the name of one firewall filter to be evaluated when packets are transmitted on the interface.



NOTE: Output filters do not work for broadcast and multicast traffic, including VPLS traffic (except in MX Series routers with MPC/MIC interfaces), as shown in [“Example: Applying a Filter to an Interface” on page 23](#).



NOTE: On an MX Series router, you cannot apply as an output filter, a firewall filter configured at the [edit firewall filter family ccc] hierarchy level. Firewall filters configured for the family ccc statement can be applied only as input filters.

In the **output-list** statement, list the names of filters to evaluate when packets are transmitted on the interface. You can include up to 16 filter names.

You can use the same filter one or more times. On M Series routers (except the M320 and M120 routers), if you apply a firewall filter or policer to multiple interfaces, the filter or policer acts on the sum of traffic entering or exiting those interfaces.

On T Series, M120, and M320 routers, interfaces are distributed among multiple packet forwarding components. Therefore, on these routers, if you apply a firewall filter or policer to multiple interfaces, the filter or policer acts on the traffic stream entering or exiting each interface, regardless of the sum of traffic on the multiple interfaces.

For more information on Understanding Ethernet Frame Statistics, see the *MX Series Layer 2 Configuration Guide*.

If you apply the filter to the interface **lo0**, it is applied to packets received or transmitted by the Routing Engine. You cannot apply MPLS filters to the management interface (**fxp0** or **em0**) or the loopback interface (**lo0**).

Filters applied at the **[set interfaces lo0 unit 0 family any filter input]** hierarchy level are not installed on T4000 Type 5 FPCs.

For more information about firewall filters, see the *Routing Policy Feature Guide for Routing Devices*. For more information about MPLS filters, see the *Junos OS MPLS Applications Library for Routing Devices*.

See also the following sections:

- [Defining Interface Groups in Firewall Filters on page 22](#)
- [Filter-Based Forwarding on the Output Interface on page 22](#)
- [Example: Applying a Filter to an Interface on page 23](#)

Defining Interface Groups in Firewall Filters

When applying a firewall filter, you can define an interface to be part of an *interface group*. Packets received on that interface are tagged as being part of the group. You can then match these packets using the **interface-group** match statement, as described in the *Routing Policy Feature Guide for Routing Devices*.

To define the interface to be part of an interface group, include the **group** statement:

```
group filter-group-number;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number* family *family* filter]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family* filter]**



NOTE: The number 0 is not a valid interface group number.

Filter-Based Forwarding on the Output Interface

If port-mirrored packets are to be distributed to multiple monitoring or collection interfaces, based on patterns in packet headers, it is helpful to configure a filter-based forwarding (FBF) filter on the port-mirroring egress interface.

When an FBF filter is installed as an output filter, a packet that is forwarded to the filter has already undergone at least one route lookup. After the packet is classified at the egress interface by the FBF filter, it is redirected to another routing table for additional route lookup. To avoid packet looping inside the Packet Forwarding Engine, the route lookup in the latter routing table (designated by an FBF routing instance) must result in a different next hop from any next hop specified in a table that has already been applied to the packet.

If an input interface is configured for FBF, the source lookup is disabled for those packets headings to a different routing instance, since the routing table is not set up to handle the source lookup.

For more information about FBF configuration, see the *Junos OS Routing Protocols Library for Routing Devices*. For more information about port mirroring, see the *Junos OS Services Interfaces Library for Routing Devices*.

Example: Applying a Filter to an Interface

Input Filter for VPLS Traffic

For M Series and T Series routers only, apply an input filter to VPLS traffic. Output filters do not work for broadcast and multicast traffic, including VPLS traffic. Note that on MX Series routers with MPC/MIC interfaces, the VPLS filters on the egress is applicable to broadcast, multicast, and unknown unicast traffic.



NOTE:

```
[edit interfaces]
fe-2/2/3 {
  vlan-tagging;
  encapsulation vlan-vpls;
  unit 601 {
    encapsulation vlan-vpls;
    vlan-id 601;
    family vpls {
      filter {
        input filter1; # Works for multicast destination MAC address
        output filter1; # Does not work for multicast destination MAC address
      }
    }
  }
}
[edit firewall]
family vpls {
  filter filter1 {
    term 1 {
      from {
        destination-mac-address {
          01:00:0c:cc:cc:cd/48;
        }
      }
      then {
        discard;
      }
    }
  }
}
```

```
term 2 {  
  then {  
    accept;  
  }  
}  
}
```

Filter-Based Forwarding at the Output Interface

The following example illustrates the configuration of filter-based forwarding at the output interface. In this example, the packet flow follows this path:

1. A packet arrives at interface **fe-1/2/0.0** with source and destination addresses **10.50.200.1** and **10.50.100.1** respectively.
2. The route lookup in routing table **inet.0** points to the egress interface **so-0/0/3.0**.
3. The output filter installed at **so-0/0/3.0** redirects the packet to routing table **fbf.inet.0**.
4. The packet matches the entry **10.50.100.0/25** in the **fbf.inet.0** table, and finally leaves the router from interface **so-2/0/0.0**.

```
[edit interfaces]  
so-0/0/3 {  
  unit 0 {  
    family inet {  
      filter {  
        output fbf;  
      }  
      address 10.50.10.2/25;  
    }  
  }  
}  
fe-1/2/0 {  
  unit 0 {  
    family inet {  
      address 10.50.50.2/25;  
    }  
  }  
}  
so-2/0/0 {  
  unit 0 {  
    family inet {  
      address 10.50.20.2/25;  
    }  
  }  
}  
[edit firewall]  
filter fbf {  
  term 0 {  
    from {  
      source-address {  
        10.50.200.0/25;  
      }  
    }  
    then routing-instance fbf;  
  }  
}
```

```

term d {
  then count d;
}
[edit routing-instances]
fbf {
  instance-type forwarding;
  routing-options {
    static {
      route 10.50.100.0/25 next-hop so-2/0/0.0;
    }
  }
}
[edit routing-options]
interface-routes {
  rib-group inet fbf-group;
}
static {
  route 10.50.100.0/25 next-hop 10.50.10.1;
}
rib-groups {
  fbf-group {
    import-rib [inet.0 fbf.inet.0];
  }
}

```

Related Documentation

- *Routing Policy Feature Guide for Routing Devices*

Enabling Source Class and Destination Class Usage

For interfaces that carry IPv4, IPv6, MPLS, or peer AS billing traffic, you can maintain packet counts based on the entry and exit points for traffic passing through your network. Entry and exit points are identified by source and destination prefixes grouped into disjoint sets defined as *source classes* and *destination classes*. You can define classes based on a variety of parameters, such as routing neighbors, autonomous systems, and route filters.

Source class usage (SCU) counts packets sent to customers by performing lookup on the IP source address. SCU makes it possible to track traffic originating from specific prefixes on the provider core and destined for specific prefixes on the customer edge. You must enable SCU accounting on both the inbound and outbound physical interfaces, and the route for the source of the packet must be in located in the forwarding table.



NOTE: SCU and DCU accounting do not work with directly connected interface routes. Source class usage does not count packets coming from sources with direct routes in the forwarding table because of software architecture limitations.

Destination class usage (DCU) counts packets from customers by performing lookup of the IP destination address. DCU makes it possible to track traffic originating from the customer edge and destined for specific prefixes on the provider core router.



NOTE: SCU and DCU accounting are supported on the J Series router only for IPv4 and IPv6 traffic.




NOTE: We recommend that you stop the network traffic on an interface before you modify the DCU or SCU configuration for that interface. Modifying the DCU or SCU configuration without stopping the traffic might corrupt the DCU or SCU statistics. Before you restart the traffic after modifying the configuration, enter the `clear interfaces statistics` command.

Figure 2 on page 26 illustrates an Internet service provider (ISP) network. In this topology, you can use DCU to count packets customers send to specific prefixes. For example, you can have three counters, one per customer, that count the packets destined for prefix **210.210/16** and **220.220/16**.

You can use SCU to count packets the provider sends from specific prefixes. For example, you can count the packets sent from prefix **210.210/16** and **215.215/16** and transmitted on a specific output interface.

Figure 2: Prefix Accounting with Source and Destination Classes



You can configure up to 126 source classes and 126 destination classes. For each interface on which you enable destination class usage and source class usage, the Junos OS maintains an interface-specific counter for each corresponding class up to the 126 class limit.



NOTE: For transit packets exiting the router through the tunnel, forwarding path features, such as RPF, forwarding table filtering, source class usage, and destination class usage are not supported on the interfaces you configure as the output interface for tunnel traffic. For firewall filtering, you must allow the output tunnel packets through the firewall filter applied to input traffic on the interface that is the next-hop interface towards the tunnel destination.

**NOTE:**

Performing DCU accounting when an output service is enabled produces inconsistent behavior in the following configuration:

- Both SCU input and DCU are configured on the packet input interface.
- SCU output is configured on the packet output interface.
- Interface services is enabled on the output interface.

For an incoming packet with source and destination prefixes matching the SCU and DCU classes respectively configured in the router, both SCU and DCU counters will be incremented. This behavior is not harmful or negative. However, it is inconsistent with non-serviced packets, in that only the SCU count will be incremented (because the SCU class ID will override the DCU class ID in this case).

To enable packet counting on an interface, include the **accounting** statement:

```
accounting {
  destination-class-usage;
  source-class-usage {
    direction;
  }
}
```

direction can be one of the following:

- **input**—Configure at least one expected ingress point.
- **output**—Configure at least one expected egress point.
- **input output**—On a single interface, configure at least one expected ingress point and one expected egress point.

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family (inet | inet6 | mpls)]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family (inet | inet6 | mpls)]

For SCU to work, you must configure at least one input interface and at least one output interface.

The ability to count a single packet for both SCU and DCU accounting depends on the underlying physical interface.

- For traffic over interfaces on Junos Trio chipset-based Packet Forwarding Engine, a single incoming packet is counted for both SCU and DCU accounting if both SCU and DCU are configured. To ensure the outgoing packet is counted, include the **source-class-usage output** statements in the configuration of the outgoing interface.

- For traffic over other interfaces, an incoming packet is counted only once, and SCU takes priority over DCU. This means that when a packet arrives on an interface on which you include the **source-class-usage input** and **destination-class-usage** statements in the configuration, and when the source and destination both match accounting prefixes, the Junos OS associates the packet with the source class only.

For traffic over interfaces on Junos Trio chipset-based Packet Forwarding Engine, SCU and DCU accounting is performed after output filters are evaluated. If a packet matches a firewall filter match condition, the packet is included in SCU or DCU accounting except in the case where the action of the matched term is **discard**.

On T Series, M120, and M320 routers, the source class and destination classes are not carried across the router fabric. The implications of this are as follows:

- On T Series, M120, and M320 routers, SCU and DCU accounting is performed before the packet enters the fabric.
- On M7i, M10i, M120, and M320 routers, on MX Series routers with non-Trio chipset-based Packet Forwarding Engine, and on T Series routers, SCU and DCU accounting is performed before output filters are evaluated. Consequently, if a packet matches a firewall filter match condition, the packet is included in SCU or DCU accounting; the packet is counted for any term action (including the **discard** action).
- On M120, M320, and T Series routers, the **destination-class** and **source-class** statements are supported at the **[edit firewall family *family-name* filter *filter-name* term *term-name* from]** hierarchy level only for the filter applied to the forwarding table. On M7i, M10i, and MX Series routers, these statements are supported.

Once you enable accounting on an interface, the Junos OS maintains packet counters for that interface, with separate counters for **inet**, **inet6**, and **mpls** protocol families. You must then configure the source class and destination class attributes in policy action statements, which must be included in forwarding-table export policies.



NOTE: When configuring policy action statements, you can configure only one source class for each matching route. In other words, more than one source class cannot be applied to the same route.

In Junos OS Release 9.3 and later, you can configure SCU accounting for Layer 3 VPNs configured with the **vrf-table-label** statement. Include the **source-class-usage** statement at the **[edit routing-instances *routing-instance-name* vrf-table-label]** hierarchy level. The **source-class-usage** statement at this hierarchy level is supported only for the virtual routing and forwarding (VRF) instance type.



NOTE: DCU counters cannot be enabled on the label-switched interface (LSI) that is created dynamically when the **vrf-table-label** statement is configured within a VRF. For more information, see the *Junos OS VPNs Library for Routing Devices*.

For a complete discussion about source and destination class accounting profiles, see the *Network Management Administration Guide for Routing Devices*. For more information about MPLS, see the *Junos OS MPLS Applications Library for Routing Devices*.

Examples: Enabling Source Class and Destination Class Usage

Configure DCU and SCU output on one interface:

```
[edit]
interfaces {
  so-6/1/0 {
    unit 0 {
      family inet {
        accounting {
          destination-class-usage;
          source-class-usage {
            output;
          }
        }
      }
    }
  }
}
```

Complete SCU Configuration

Source routers A and B use loopback addresses as the prefixes to be monitored. Most of the configuration tasks and actual monitoring occur on transit Router SCU.

The loopback address on Router A contains the origin of the prefix that is to be assigned to source class A on Router SCU. However, no SCU processing happens on this router. Therefore, configure Router A for basic OSPF routing and include your loopback interface and interface **so-0/0/2** in the OSPF process.

Router A

```
[edit]
interfaces {
  so-0/0/2 {
    unit 0 {
      family inet {
        address 10.255.50.2/24;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.255.192.10/32;
      }
    }
  }
}
protocols {
  ospf {
    area 0.0.0.0 {
      interface so-0/0/2.0;
      interface lo0.0;
    }
  }
}
```

```
}
```

Router SCU Last, apply the policy to the forwarding table.

Router SCU handles the bulk of the activity in this example. On Router SCU, enable source class usage on the inbound and outbound interfaces at the **[edit interfaces interface-name unit unit-number family inet accounting]** hierarchy level. Make sure you specify the expected traffic: input, output, or, in this case, both.

Next, configure a route filter policy statement that matches the prefixes of the loopback addresses from routers A and B. Include statements in the policy that classify packets from Router A in one group named **scu-class-a** and packets from Router B in a second class named **scu-class-b**. Notice the efficient use of a single policy containing multiple terms.

```
[edit]
interfaces {
  so-0/0/1 {
    unit 0 {
      family inet {
        accounting {
          source-class-usage {
            input;
            output;
          }
        }
        address 10.255.50.1/24;
      }
    }
  }
  so-0/0/3 {
    unit 0 {
      family inet {
        accounting {
          source-class-usage {
            input;
            output;
          }
        }
        address 10.255.10.3/24;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.255.6.111/32;
      }
    }
  }
}
protocols {
  ospf {
    area 0.0.0.0 {
      interface so-0/0/1.0;
      interface so-0/0/3.0;
```



```

    }
  }
}
routing-options {
  forwarding-table {
    export scu-policy;
  }
}
policy-options {
  policy-statement scu-policy {
    term 0 {
      from {
        route-filter 10.255.192.0/24 orlonger;
      }
      then source-class scu-class-a;
    }
    term 1 {
      from {
        route-filter 10.255.165.0/24 orlonger;
      }
      then source-class scu-class-b;
    }
  }
}
}

```

Router B Just as Router A provides a source prefix, Router B's loopback address matches the prefix assigned to **scu-class-b** on Router SCU. Again, no SCU processing happens on this router, so configure Router B for basic OSPF routing and include your loopback interface and interface **so-0/0/4** in the OSPF process.

```

interfaces {
  so-0/0/4 {
    unit 0 {
      family inet {
        address 10.255.10.4/24;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.255.165.226/32;
      }
    }
  }
}
protocols {
  ospf {
    area 0.0.0.0 {
      interface so-0/0/4.0;
      interface lo0.0;
    }
  }
}
}

```

**Enabling Packet
Counting for Layer 3
VPNs**

You can use SCU and DCU to count packets on Layer 3 VPNs. To enable packet counting for Layer 3 VPN implementations at the egress point of the MPLS tunnel, you must configure a virtual loopback tunnel interface (**vt**) on the PE router, map the virtual routing and forwarding (VRF) instance type to the virtual loopback tunnel interface, and send the traffic received from the VPN out the source class output interface, as shown in the following example:

Configure a virtual loopback tunnel interface on a provider edge router equipped with a tunnel PIC:

```
[edit interfaces]
vt-0/3/0 {
  unit 0 {
    family inet {
      accounting {
        source-class-usage {
          input;
        }
      }
    }
  }
}
```

Map the VRF instance type to the virtual loopback tunnel interface.

In Junos OS Release 9.3 and later, you can configure SCU accounting for Layer 3 VPNs configured with the **vrf-table-label** statement. Include the **source-class-usage** statement at the **[edit routing-instances routing-instance-name vrf-table-label]** hierarchy level. The **source-class-usage** statement at this hierarchy level is supported only for the virtual routing and forwarding (VRF) instance type. DCU is not supported when the **vrf-table-label** statement is configured. For more information, see the *Junos OS VPNs Library for Routing Devices*.

```
[edit routing-instances]
VPN-A {
  instance-type vrf;
  interface at-2/1/1.0;
  interface vt-0/3/0.0;
  route-distinguisher 10.255.14.225:100;
  vrf-import import-policy-A;
  vrf-export export-policy-A;
  protocols {
    bgp {
      group to-r4 {
        local-address 10.27.253.1;
        peer-as 400;
        neighbor 10.27.253.2;
      }
    }
  }
}
```

Send traffic received from the VPN out the source class output interface:

```
[edit interfaces]
at-2/1/0 {
  unit 0 {
```

```

family inet {
  accounting {
    source-class-usage {
      output;
    }
  }
}

```

For more information about VPNs, see the *Junos OS VPNs Library for Routing Devices*. For more information about virtual loopback tunnel interfaces, see the *Junos OS Services Interfaces Library for Routing Devices*.

Related Documentation

- [accounting on page 105](#)
- [destination-classes](#)
- [family on page 127](#)
- [forward-and-send-to-re on page 133](#)
- [source-classes](#)
- [targeted-broadcast on page 170](#)
- [unit on page 176](#)

Configuring the Protocol Family

For each logical interface, you can configure one or more of the following protocols that run on the interface:

- **any**—Protocol-independent family used for Layer 2 packet filtering. This option is not supported on J Series routers and on T4000 Type 5 FPCs.
- **bridge**—(M Series and T Series routers only) Configure only when the physical interface is configured with **ethernet-bridge** type encapsulation or when the logical interface is configured with **vlan-bridge** type encapsulation. You can optionally configure this protocol family for the logical interface on which you configure VPLS.
- **ccc**—Circuit cross-connect (CCC). You can configure this protocol family for the logical interface of CCC physical interfaces. When you use this encapsulation type, you can configure the **ccc** family only.
- **inet**—IP. You must configure this protocol family for the logical interface to support IP protocol traffic, including Open Shortest Path First (OSPF), Border Gateway Protocol (BGP), Internet Control Message Protocol (ICMP), and Internet Protocol Control Protocol (IPCP).
- **inet6**—IP version 6 (IPv6). You must configure this protocol family for the logical interface to support IPv6 protocol traffic, including Routing Information Protocol for IPv6 (RIPng), Intermediate System-to-Intermediate System (IS-IS), BGP, and Virtual Router Redundancy Protocol for IPv6 (VRRP). For more information about IPv6, see [“IPv6 Overview” on page 35](#).

- **iso**—International Organization for Standardization (ISO). You must configure this protocol family for the logical interface to support IS-IS traffic.
- **mlfr-uni-nni**—Multilink Frame Relay (MLFR) FRF.16 user-to-network network-to-network (UNI NNI). You must configure this protocol or **mlfr-end-to-end** for the logical interface to support link services and voice services bundling.
- **mlfr-end-to-end**—Multilink Frame Relay end-to-end. You must configure this protocol or multilink Point-to-Point Protocol (MLPPP) for the logical interface to support multilink bundling.
- **mlppp**—MLPPP. You must configure this protocol (or **mlfr-end-to-end**) for the logical interface to support multilink bundling.
- **mpls**—Multiprotocol Label Switching (MPLS). You must configure this protocol family for the logical interface to participate in an MPLS path.
- **tcc**—Translational cross-connect (TCC). You can configure this protocol family for the logical interface of TCC physical interfaces.
- **tnp**—Trivial Network Protocol. This protocol is used to communicate between the Routing Engine and the router's packet forwarding components. The Junos OS automatically configures this protocol family on the router's internal interfaces only, as discussed in *Understanding Internal Ethernet Interfaces*.
- **vpls**—M Series and T Series routers support Virtual Private LAN service (VPLS). You can optionally configure this protocol family for the logical interface on which you configure VPLS. VPLS provides an Ethernet-based point-to-multipoint Layer 2 VPN to connect customer edge (CE) routers across an MPLS backbone. When you configure a VPLS encapsulation type, the **family vpls** statement is assumed by default.

MX Series routers support dynamic profiles for VPLS pseudowires, VLAN identifier translation, and automatic bridge domain configuration.

For more information about VPLS, see the *Junos OS VPNs Library for Routing Devices* and the *Junos OS, Release 14.1*.

To configure the logical interface's protocol family, include the **family** statement, specifying the selected family. To configure more than one protocol **family** on a logical interface, include multiple **family** statements. Following is the minimum configuration:

```
family family {  
  mtu size;  
  multicast-only;  
  no-redirects;  
  primary;  
  address address {  
    destination address;  
    broadcast address;  
    preferred;  
    primary;  
  }  
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

IPv6 Overview

IP version 4 (IPv4) has been widely deployed and used to network the Internet today. With the rapid growth of the Internet, enhancements to IPv4 are needed to support the influx of new subscribers, Internet-enabled devices, and applications. IPv6 is designed to enable the global expansion of the Internet.

IPv6 builds upon the functionality of IPv4, providing improvements to addressing, configuration and maintenance, and security.

IPv6 is defined in the following documents:

- RFC 2373, *IP Version 6 Addressing Architecture*
- RFC 2460, *Internet Protocol, Version 6 (IPv6)*

IPv4-to-IPv6 Transition

Implementing IPv6 requires a transition mechanism to allow interoperability between IPv6 nodes (both routers and hosts) and IPv4 nodes. The transition mechanism is the key factor in the successful deployment of IPv6. Because millions of IPv4 nodes already exist, upgrading every node to IPv6 at the same time is not feasible.

As a result, transition from IPv4 to IPv6 happens gradually, allowing nodes to be upgraded independently and without disruption to other nodes. While a gradual upgrade occurs, compatibility between IPv6 and IPv4 nodes becomes a requirement. Otherwise, an IPv6 node would not be able to communicate with an IPv4 node.

Transition mechanisms allow IPv6 and IPv4 nodes to coexist together in the same network, and make gradual upgrading possible. The transition mechanism supported by the Junos OS is tunneling. Tunnels allow IPv6 packets to be encapsulated into IPv4 headers and sent across an IPv4 infrastructure. For more information about configuring tunnels to support IPv4-to-IPv6 transition, see the *Junos OS Services Interfaces Library for Routing Devices*.

VRRP Properties

The Virtual Router Redundancy Protocol (VRRP) provides a much faster switchover to a backup router when the default router fails. Using VRRP, a backup router can take over a failed default router within a few seconds. This is done with minimum amount of VRRP traffic and without any interactions with the hosts.

For more information on VRRP properties, see the *Junos OS High Availability Library for Routing Devices*.

Related Documentation • [Understanding Internal Ethernet Interfaces](#)

Configuring the Interface Address

You assign an address to an interface by specifying the address when configuring the protocol family. For the **inet** or **inet6** family, configure the interface IP address. For the **iso** family, configure one or more addresses for the loopback interface. For the **ccc**, **ethernet-switching**, **tcc**, **mpls**, **tnp**, and **vpls** families, you never configure an address.



NOTE: The point-to-point (PPP) address is taken from the loopback interface address that has the primary attribute. When the loopback interface is configured as an unnumbered interface, it takes the primary address from the donor interface.

To assign an address to an interface, include the **address** statement:

```
address address {  
    broadcast address;  
    destination address;  
    destination-profile name;  
    eui-64;  
    preferred;  
    primary;  
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family *family*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*]

In the **address** statement, specify the network address of the interface.

For each address, you can optionally configure one or more of the following:

- Broadcast address for the interface subnet—Specify this in the **broadcast** statement; this applies only to Ethernet interfaces, such as the management interface **fxp0**, **em0**, or **me0** the Fast Ethernet interface, and the Gigabit Ethernet interface.
- Address of the remote side of the connection (for point-to-point interfaces only)—Specify this in the **destination** statement.
- PPP properties to the remote end—Specify this in the **destination-profile** statement. You define the profile at the [edit access group-profile *name* **ppp**] hierarchy level (for point-to-point interfaces only).
- Whether the router or switch automatically generates the host number portion of interface addresses—The **eui-64** statement applies only to interfaces that carry IPv6 traffic, in which the prefix length of the address is 64 bits or less, and the low-order 64 bits of the address are zero. This option does not apply to the loopback interface (**lo0**)

because IPv6 addresses configured on the loopback interface must have a 128-bit prefix length.

- Whether this address is the preferred address—Each subnet on an interface has a preferred local address. If you configure more than one address on the same subnet, the preferred local address is chosen by default as the source address when you originate packets to destinations on the subnet.

By default, the preferred address is the lowest-numbered address on the subnet. To override the default and explicitly configure the preferred address, include the **preferred** statement when configuring the address.

- Whether this address is the primary address—Each interface has a primary local address. If an interface has more than one address, the primary local address is used by default as the source address when you send packets from an interface where the destination provides no information about the subnet (for example, some **ping** commands).

By default, the primary address on an interface is the lowest-numbered non-127 (in other words, non-loopback) preferred address on the interface. To override the default and explicitly configure the preferred address, include the **primary** statement when configuring the address.



NOTE: If you configure a duplicate IP address on an interface, even when the earlier interface with that IP address is disabled, a Warning message is added to the syslog and not displayed on the screen. Do not configure the same IP address of a disabled interface on another interface.

- [Configuring Interface IPv4 Addresses on page 37](#)
- [Configuring Interface IPv6 Addresses on page 39](#)

Configuring Interface IPv4 Addresses

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

To configure an IPv4 address on routers and switches running Junos OS, use the **edit interface *interface-name* unit *number* family inet address *a.b.c.d/nn*** statement at the **[edit interfaces]** hierarchy level.



NOTE: Juniper Networks routers and switches support /31 destination prefixes when used in point-to-point Ethernet configurations; however, they are not supported by many other devices, such as hosts, hubs, routers, or switches. You must determine if the peer system also supports /31 destination prefixes before configuration.

Operational Behavior of Interfaces when the Same IPv4 Address is Assigned to Them

You can configure the same IPv4 address on multiple physical interfaces. When you assign the same IPv4 address to multiple physical interfaces, the operational behavior of those interfaces differs, depending on whether they are implicitly or explicitly point-to-point.



NOTE: By default, all interfaces are assumed to be point-to-point (PPP) interfaces. For all interfaces except aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet, you can explicitly configure an interface to be a point-to-point connection.

The following examples show the sample configuration of assigning the same IPv4 address to implicitly and explicitly point-to-point interfaces, and their corresponding **show interfaces terse** command outputs to see their operational status.

Configuring same IPv4 address on implicitly PPP interfaces:

```
[edit]
user@host# show
ge-0/1/0 {
  unit 0 {
    family inet {
      address 200.1.1.1/24;
    }
  }
}

ge-3/0/1 {
  unit 0 {
    family inet {
      address 200.1.1.1/24;
    }
  }
}
```

The sample output shown below for the above configuration reveals that only **ge-0/1/0.0** was assigned the same IPv4 address **200.1.1.1/24** and its **link** state was **up**, while **ge-3/0/1.0** was not assigned the IPv4 address, though its **link** state was **up**, which means that it will be operational only when it gets a unique IPv4 address other than **200.1.1.1/24**.

```
user@host> show interfaces terse ge*
Interface           Admin Link Proto  Local Remote
ge-0/1/0             up   up   up
ge-0/1/0.0           up   up   inet   200.1.1.1/24
                    multiservice
ge-0/1/1             up   down
ge-3/0/0             up   down
ge-3/0/1             up   up
ge-3/0/1.0           up   up   inet   multiservice
```


Configuring same IPv4 address on explicitly PPP interfaces:

```
[edit]
user@host# show
so-0/0/0 {
  unit 0 {
    family inet {
      address 200.1.1.1/24;
    }
  }
}
so-0/0/3 {
  unit 0 {
    family inet {
      address 200.1.1.1/24;
    }
  }
}
```

The sample output shown below for the above configuration reveals that both **so-0/0/0.0** and **so-0/0/3.0** were assigned the same IPv4 address **200.1.1.1/24** and that their link states were down, which means that to make them operational at least one of them will have to be configured with a unique IPv4 address other than **200.1.1.1/24**.

```
user@host> show interfaces terse so*
Interface      Admin Link Proto  Local          Remote
so-0/0/0        up   up
so-0/0/0.0      up   down inet    200.1.1.1/24
so-0/0/1        up   up
so-0/0/2        up   down
so-0/0/3        up   up
so-0/0/3.0      up   down inet    200.1.1.1/24
so-1/1/0        up   down
so-1/1/1        up   down
so-1/1/2        up   up
so-1/1/3        up   up
so-2/0/0        up   up
so-2/0/1        up   up
so-2/0/2        up   up
so-2/0/3        up   down
```

Configuring Interface IPv6 Addresses

NOTE: IPv6 is not currently supported for the QFX Series.

You represent IP version 6 (IPv6) addresses in hexadecimal notation using a colon-separated list of 16-bit values.

You assign a 128-bit IPv6 address to an interface by including the **address** statement:

```
address aaaa:bbbb:...:zzzz/nn;
```



NOTE: You cannot configure a subnet zero IPv6 address because RFC 2461 reserves the subnet-zero address for anycast addresses, and Junos OS complies with the RFC.

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet6]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet6]

The double colon (::) represents all bits set to 0, as shown in the following example:

```
interfaces fe-0/0/1 {
  unit 0 {
    family inet6 {
      address fec0:1:1::2/64;
    }
  }
}
```



NOTE: You must manually configure the router or switch advertisement and advertise the default prefix for autoconfiguration to work on a specific interface.

**Related
Documentation**

- [Configuring IPCP Options on page 42](#)
- [Configuring Default, Primary, and Preferred Addresses and Interfaces on page 50](#)

Configuring ICCP for MC-LAG

For multichassis link aggregation (MC-LAG), you must configure Inter-Control Center Communications Protocol (ICCP) to exchange information between two MC-LAG peers.

To enable ICCP, include the **iccp** statement at the **[edit protocols]** hierarchy level:

```
[edit protocols]
iccp {
  authentication-key string;
  local-ip-addr ipv4-address;
  peer ip-address {
    authentication-key string;
    liveness-detection {
      detection-time {
        threshold milliseconds;
      }
    }
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
  }
}
```

```

    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (1 | automatic);
}
local-ip-addr ipv4-address;
redundancy-group-id-list [ redundancy-groups ];
session-establishment-hold-time value;
}
session-establishment-hold-time value;
traceoptions;
}

```

The **local-ip-address** statement sets the source address. This could be a specified address or interface address. The **session-establishment-hold-time** statement determines whether a chassis takes over as the master at the ICCP session.

The **authentication-key** statement is provided by TCP Message Digest 5 (md5) option for an ICCP TCP session. The **redundancy-group-id-list** statement specifies the redundancy groups between ICCP peers and the **liveness-detection** hierarchy configures Bidirectional Forwarding Detection (BFD) protocol options.



NOTE: ICCP is based on TCP and it uses IP routes to reach the MC-LAG peer. To ensure that the ICCP session is as resilient as possible, we recommend that you configure alternative routes between the ICCP end-point IP addresses. Alternatively, configure a LAG interface that has two or more interfaces between the MC-LAG pairs to prevent session failure when there are no alternative routes.

For Inter-Control Center Communications Protocol (ICCP) in a multichassis link aggregation group (MC-LAG) configured in an active-active bridge domain, you must ensure that you configure the same peer IP address hosting the MC-LAG by including the **peer ip-address** statement at the **[edit protocols iccp]** hierarchy level and the **multi-chassis-protection peer ip-address** statement at the **[edit interfaces interface-name]** hierarchy level. Multichassis protection reduces the configuration at the logical interface level for MX Series routers with multichassis aggregated Ethernet (MC-AE) interfaces. If the ICCP is UP and the interchassis data link (ICL) comes UP, the router configured as standby will bring up the MC-AE interfaces shared with the peer active-active node specified by the **peer** statement.

For example, the following statements illustrate how the same peer IP address can be configured for both the ICCP peer and multichassis protection link:

```

set interfaces ae1 unit 0 multi-chassis-protection 10.255.34.112 interface ae0.0
set protocols iccp peer 10.255.34.112 redundancy-group-id-list 1

```

Although you can commit an MC-LAG configuration with various parameters defined for it, you can configure multichassis protection between two peers without configuring the ICCP peer address. You can also configure multiple ICCP peers and commit such a configuration.

Related • [\[edit protocols\] Hierarchy Level](#)
Documentation

Configuring IPCP Options

For interfaces with PPP encapsulation, you can configure IPCP to negotiate IP address assignments and to pass network-related information such as Windows Name Service (WINS) and Domain Name System (DNS) servers, as defined in RFC 1877, *PPP Internet Protocol Control Protocol Extensions for Name Server Addresses*.



NOTE: The Junos OS does not request name servers from the remote end; the software does, however, send name servers to the remote end if requested.

On the logical interface, the following PPP encapsulation types are supported:

- `atm-mlppp-llc`
- `atm-ppp-llc`
- `atm-ppp-vc-mux`
- `multilink-ppp`

When you enable a PPP interface, you can configure an IP address, enable the interface to negotiate an IP address assignment from the remote end, or allow the interface to be unnumbered. You can also assign a destination profile to the remote end. The destination profile includes PPP properties, such as primary and secondary DNS and NetBIOS Name Servers (NBNSs). These options are described in the following sections:

- [Configuring an IP Address for an Interface on page 42](#)
- [Negotiating an IP Address Assignment from the Remote End on page 43](#)
- [Configuring an Interface to Be Unnumbered on page 43](#)
- [Assigning a Destination Profile to the Remote End on page 43](#)

Configuring an IP Address for an Interface

You can configure an IP address for the interface by including the **address** statement in the configuration. For more information, see [“Configuring the Interface Address” on page 36](#).

If you include the **address** statement in the configuration, you cannot include the **negotiate-address** or **unnumbered-address** statement in the configuration.

When you include the **address** statement in the interface configuration, you can assign PPP properties to the remote end, as shown in [“Assigning a Destination Profile to the Remote End” on page 43](#).



NOTE: The option to negotiate an IP address is not allowed in MLFR and MFR encapsulations.

Negotiating an IP Address Assignment from the Remote End

To enable the interface to obtain an IP address from the remote end, include the **negotiate-address** statement:

```
negotiate-address;
```

You can include this statement at the following hierarchy levels:

- [edit **interfaces** *interface-name* **unit** *logical-unit-number* **family** *inet*]
- [edit **logical-systems** *logical-system-name* **interfaces** *interface-name* **unit** *logical-unit-number* **family** *inet*]

If you include the **negotiate-address** statement in the configuration, you cannot include the **address** or **unnumbered-address** statement in the configuration.

Configuring an Interface to Be Unnumbered

To configure an interface to be unnumbered, include the **unnumbered-address** and **destination** statements in the configuration:

```
unnumbered-address interface-name destination address;
```

The **unnumbered-address** statement enables the local address to be derived from the specified interface. The interface name must include a logical unit number and must have a configured address (see “[Configuring the Interface Address](#)” on page 36). Specify the IP address of the remote interface with the **destination** statement.

You can include these statements at the following hierarchy levels:

- [edit **interfaces** *interface-name* **unit** *logical-unit-number* **family** *inet*]
- [edit **logical-systems** *logical-system-name* **interfaces** *interface-name* **unit** *logical-unit-number* **family** *inet*]

If you include the **unnumbered-address** statement in the configuration, you cannot include the **address** or **negotiate-address** statement in the interface configuration.

When you include the **unnumbered-address** statement in the interface configuration, you can assign PPP properties to the remote end, as shown in “[Assigning a Destination Profile to the Remote End](#)” on page 43.

Assigning a Destination Profile to the Remote End

When you include the **address** or **unnumbered-address** statement in the interface configuration, you can assign PPP properties to the remote end. To do this, include the **destination-profile** statement:

```
destination-profile name;
```

You can include this statement at the following hierarchy levels:

- [edit [interfaces](#) *interface-name* [unit](#) *logical-unit-number* [family](#) [inet](#) [address](#) *address*]
- [edit [interfaces](#) *interface-name* [unit](#) *logical-unit-number* [family](#) [inet](#) [unnumbered-address](#) *interface-name*]
- [edit [logical-systems](#) *logical-system-name* [interfaces](#) *interface-name* [unit](#) *logical-unit-number* [family](#) [inet](#) [address](#) *address*]
- [edit [logical-systems](#) *logical-system-name* [interfaces](#) *interface-name* [unit](#) *logical-unit-number* [family](#) [inet](#) [unnumbered-address](#) *interface-name*]

The profile name is a PPP group profile. You define the profile by including the following statements at the [edit [access group-profile](#) *name* [ppp](#)] hierarchy level:

```
[edit access group-profile name ppp]
framed-pool pool-id;
interface-id interface-id;
primary-dns primary-dns;
primary-wins primary-win-server;
secondary-dns secondary-dns;
secondary-wins secondary-wins;
```

For more information about PPP group profiles, see the *Junos OS Administration Library for Routing Devices*.

**Related
Documentation**

- [Configuring the Interface Address on page 36](#)
- *Junos OS Administration Library for Routing Devices*

Configuring an Unnumbered Interface

When you need to conserve IP addresses, you can configure unnumbered interfaces. Setting up an unnumbered interface enables IP processing on the interface without assigning an explicit IP address to the interface. For IPv6, in which conserving addresses is not a major concern, you can configure unnumbered interfaces to share the same subnet across multiple interfaces. IPv6 unnumbered interfaces are only supported on Ethernet interfaces. The statements you use to configure an unnumbered interface depend on the type of interface you are configuring: a point-to-point interface or an Ethernet interface:

- [Configuring an Unnumbered Point-to-Point Interface on page 44](#)
- [Configuring an Unnumbered Ethernet or Demux Interface on page 45](#)

Configuring an Unnumbered Point-to-Point Interface

To configure an unnumbered point-to-point interface, configure the protocol family, but do not include the **address** statement:

```
family family;
```

You can include this statement at the following hierarchy levels:

- [edit **interfaces** *interface-name* unit *logical-unit-number*]
- [edit **logical-systems** *logical-system-name* **interfaces** *interface-name* unit *logical-unit-number*]



NOTE: For interfaces with PPP encapsulation, you can configure an unnumbered interface by including the `unnumbered-interface` statement in the configuration. For more information, see [“Configuring IPCP Options” on page 42](#).

When configuring unnumbered interfaces, you must ensure that a source address is configured on some interface in the router. This address is the default address. We recommend that you do this by assigning an address to the loopback interface (**lo0**), as described in *Configuring the Loopback Interface*. If you configure an address (other than a martian) on the **lo0** interface, that address is always the default address, which is preferable because the loopback interface is independent of any physical interfaces and therefore is always accessible.

Example: Configuring an Unnumbered Point-to-Point Interface

Configure an unnumbered point-to-point interface:

```
[edit]
interfaces {
  so-6/1/0 {
    unit 0 {
      family inet;
      family iso;
    }
  }
}
```

Configuring an Unnumbered Ethernet or Demux Interface

To configure an unnumbered Ethernet or demultiplexing interface, include the `unnumbered-address` statement in the configuration:

```
unnumbered-address interface-name;
```

You can include this statement at the following hierarchy levels:

- [edit **interfaces** *interface-name* unit *logical-unit-number* **family** *family*]
- [edit **logical-systems** *logical-system-name* **interfaces** *interface-name* unit *logical-unit-number* **family** *family*]

For dynamic profiles, include the `unnumbered-address` statement at the following hierarchy levels:

- [edit **dynamic-profiles** *profile-name* **interfaces** *interface-name* unit *logical-unit-number* **family** *family*]

- `[edit dynamic-profiles profile-name interfaces demux0 unit logical-unit-number family family]`

The **unnumbered-address** statement currently supports configuration of unnumbered demux interfaces only for the IPv4 address family. You can configure unnumbered Ethernet interfaces for both IPv4 and IPv6 address families.

The interface that you configure to be unnumbered *borrow*s an assigned IP address from another interface, and is referred to as the *borrower interface*. The interface from which the IP address is borrowed is referred to as the *donor interface*. In the **unnumbered-address** statement, **interface-name** specifies the donor interface. For an unnumbered Ethernet interface, the donor interface can be an Ethernet, ATM, SONET, or loopback interface that has a logical unit number and configured IP address and is not itself an unnumbered interface. For an unnumbered IP demultiplexing interface, the donor interface can be an Ethernet or loopback interface that has a logical unit number and configured IP address and is not itself an unnumbered interface. In addition, for either Ethernet or demux, the donor interface and the borrower interface must be members of the same routing instance and the same logical system.

When you configure an unnumbered Ethernet or demux interface, the IP address of the donor interface becomes the source address in packets generated by the unnumbered interface.

You can configure a host route that points to an unnumbered Ethernet or demux interface. For information about host routes, see the *Junos OS MPLS Applications Library for Routing Devices*.

For more information, see the following sections:

- [Configuring a Preferred Source Address for Unnumbered Ethernet or Demux Interfaces on page 46](#)
- [Configuring Static Routes on Unnumbered Ethernet Interfaces on page 47](#)
- [Restrictions for Configuring Unnumbered Ethernet Interfaces on page 48](#)
- [Example: Configuring an Unnumbered Ethernet Interface on page 48](#)
- [Example: Configuring the Preferred Source Address for an Unnumbered Ethernet Interface on page 49](#)
- [Example: Configuring an Unnumbered Ethernet Interface as the Next Hop for a Static Route on page 49](#)

For additional information about dynamic-profiles, see *Dynamic Profiles Overview*.

Configuring a Preferred Source Address for Unnumbered Ethernet or Demux Interfaces

When a loopback interface with multiple secondary IP addresses is configured as the donor interface for an unnumbered Ethernet or demux interface, you can optionally specify any one of the loopback interface's secondary addresses as the preferred source address for the unnumbered Ethernet or demux interface. This feature enables you to use an IP address other than the primary IP address on some of the unnumbered Ethernet or demux interfaces in your network.

To configure a secondary address on a loopback donor interface as the preferred source address for an unnumbered Ethernet or demux interface, include the **preferred-source-address** option in the **unnumbered-address** statement:

```
unnumbered-address interface-name <preferred-source-address address>;
```

You can include this statement at the following hierarchy levels:

- [edit **interfaces** *interface-name* **unit** *logical-unit-number* **family** *family*]
- [edit **logical-systems** *logical-system-name* **interfaces** *interface-name* **unit** *logical-unit-number* **family** *family*]
- [edit **dynamic-profiles** *profile-name* **interfaces** *interface-name* **unit** *logical-unit-number* **family** *family*]
- [edit **dynamic-profiles** *profile-name* **interfaces** **demux0** **unit** *logical-unit-number* **family** *family*]

The following considerations apply when you configure a preferred source address on an unnumbered Ethernet or demux interface:

- The **unnumbered-address** statement currently supports the configuration of a preferred source address only for the IPv4 address family for demux interfaces, and for IPv4 and IPv6 address families for Ethernet interfaces.
- If you do not specify the preferred source address, the router uses the default primary IP address of the donor interface.
- You cannot delete an address on a donor loopback interface while it is being used as the preferred source address for an unnumbered Ethernet or demux interface.

For a configuration example that illustrates this feature, see “[Example: Configuring the Preferred Source Address for an Unnumbered Ethernet Interface](#)” on page 49.

To display the preferred source address for an unnumbered Ethernet or demux interface, use the **show interfaces** operational mode command. For information about using this command, see the [CLI Explorer](#).

Configuring Static Routes on Unnumbered Ethernet Interfaces

You can configure static routes on an unnumbered Ethernet interface. To do so, you use the **qualified-next-hop** statement to specify the unnumbered Ethernet interface as the next-hop interface for a configured static route. This feature enables you to specify independent preferences and metrics for static routes on a next-hop basis.

For a configuration example that illustrates this feature, see “[Example: Configuring an Unnumbered Ethernet Interface as the Next Hop for a Static Route](#)” on page 49.

For information about how to specify an independent preference for a static route, see the *Junos OS Routing Protocols Library for Routing Devices*.

Restrictions for Configuring Unnumbered Ethernet Interfaces

The following restrictions apply when you configure unnumbered Ethernet interfaces:

- The **unnumbered-address** statement currently supports the configuration of unnumbered Ethernet interfaces for IPv4 and IPv6 address families.
- You cannot assign an IP address to an Ethernet interface that is already configured as an unnumbered interface.
- The donor interface for an unnumbered Ethernet interface must have one or more configured IP addresses.
- The donor interface for an unnumbered Ethernet interface cannot be configured as unnumbered.
- An unnumbered Ethernet interface does not support configuration of the following **address** statement options: **arp**, **broadcast**, **primary**, **preferred**, and **vrp-group**. For information about these options, see [“Configuring the Interface Address” on page 36](#).
- Running IGMP and PIM are supported only on unnumbered Ethernet interfaces that directly face the host and have no downstream PIM neighbors. IGMP and PIM are not supported on unnumbered Ethernet interfaces that act as upstream interfaces in a PIM topology.
- Running OSPF and IS-IS on unnumbered Ethernet interfaces is not supported. However, you can run OSPF over unnumbered Ethernet interfaces configured as a Point-to-Point connection.

Example: Configuring an Unnumbered Ethernet Interface

In this example, **ge-1/0/0** is the unnumbered interface and **ge-0/0/0** is the donor interface from which **ge-1/0/0** “borrows” an IP address.

```
interfaces {
  ge-0/0/0 {
    unit 0 {
      family inet {
        address 4.4.4.1/24;
      }
    }
  }
  ge-1/0/0 {
    unit 0 {
      family inet {
        unnumbered-address ge-0/0/0.0;
      }
    }
  }
}
```

The sample configuration that is described works correctly on M and T Series routers. For unnumbered interfaces on MX Series routers, you must additionally configure static routes on an unnumbered Ethernet interface by including the **qualified-next-hop** statement

at the **[edit routing-options static route destination-prefix]** hierarchy level to specify the unnumbered Ethernet interface as the next-hop interface for a configured static route.

Example: Configuring the Preferred Source Address for an Unnumbered Ethernet Interface

In this example, loopback interface **lo0** is the donor interface from which unnumbered Ethernet interface **ge-4/0/0** “borrows” an IP address. The example also configures one of the loopback interface’s secondary addresses, 3.3.3.1, as the preferred source address for the unnumbered Ethernet interface.

```

interfaces {
  lo0 {
    unit 0 {
      family inet {
        address 2.2.2.1/32;
        address 3.3.3.1/32;
      }
    }
  }
}
interfaces {
  ge-4/0/0 {
    unit 0 {
      family inet {
        unnumbered-address lo0.0 preferred-source-address 3.3.3.1;
      }
    }
  }
}

```

Example: Configuring an Unnumbered Ethernet Interface as the Next Hop for a Static Route

In this example, **ge-0/0/0** is the unnumbered interface and a loopback interface, **lo0**, is the donor interface from which **ge-0/0/0** “borrows” an IP address. The example also configures a static route to **7.7.7.1/32** with a next hop through unnumbered interface **ge-0/0/0.0**.

```

interfaces {
  lo0 {
    unit 0 {
      family inet {
        address 5.5.5.1/32;
        address 6.6.6.1/32;
      }
    }
  }
}
interfaces {
  ge-0/0/0 {
    unit 0 {
      family inet {
        unnumbered-address lo0.0;
      }
    }
  }
}

```

```
    }  
  }  
}  
routing-options {  
  static {  
    route 7.7.7.1/32 {  
      qualified next-hop ge-0/0/0.0;  
    }  
  }  
}
```

- Related Documentation**
- [Configuring the Interface Address on page 36](#)
 - *Junos OS Administration Library for Routing Devices*

Configuring Default, Primary, and Preferred Addresses and Interfaces

The router has a default address and a primary interface, and interfaces have primary and preferred addresses.

The *default address* of the router is used as the source address on unnumbered interfaces. The routing protocol process tries to pick the default address as the router ID, which is used by protocols, including OSPF and internal BGP (IBGP).

The *primary interface* for the router is the interface that packets go out when no interface name is specified and when the destination address does not imply a particular outgoing interface.

An interface's *primary address* is used by default as the local address for broadcast and multicast packets sourced locally and sent out the interface. An interface's *preferred address* is the default local address used for packets sourced by the local router to destinations on the subnet.

The default address of the router is chosen using the following sequence:

1. The primary address on the loopback interface **lo0** that is not **127.0.0.1** is used.
2. The primary address on the primary interface is used.

To configure these addresses and interfaces, you can do the following:

- [Configuring the Primary Interface for the Router on page 50](#)
- [Configuring the Primary Address for an Interface on page 51](#)
- [Configuring the Preferred Address for an Interface on page 51](#)

Configuring the Primary Interface for the Router

The *primary interface* for the router has the following characteristics:

- It is the interface that packets go out when you type a command such as ping 255.255.255.255—that is, a command that does not include an interface name (there is no interface **type-0/0/0.0** qualifier) and where the destination address does not imply any particular outgoing interface.

- It is the interface on which multicast applications running locally on the router, such as Session Announcement Protocol (SAP), do group joins by default.
- It is the interface from which the default local address is derived for packets sourced out an unnumbered interface if there are no non-127 addresses configured on the loopback interface, lo0.

By default, the multicast-capable interface with the lowest-index address is chosen as the primary interface. If there is no such interface, the point-to-point interface with the lowest index address is chosen. Otherwise, any interface with an address could be picked. In practice, this means that, on the router, the **fxp0** or **em0** interface is picked by default.

To configure a different interface to be the primary interface, include the **primary** statement:

```
primary;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family *family*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*]

Configuring the Primary Address for an Interface

The *primary address* on an interface is the address that is used by default as the local address for broadcast and multicast packets sourced locally and sent out the interface. For example, the local address in the packets sent by a **ping interface so-0/0/0.0 255.255.255.255** command is the primary address on interface **so-0/0/0.0**. The primary address flag also can be useful for selecting the local address used for packets sent out unnumbered interfaces when multiple non-127 addresses are configured on the loopback interface, **lo0**. By default, the primary address on an interface is selected as the numerically lowest local address configured on the interface.

To set a different primary address, include the **primary** statement:

```
primary;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family *family* address *address*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family* address *address*]

Configuring the Preferred Address for an Interface

The *preferred address* on an interface is the default local address used for packets sourced by the local router to destinations on the subnet. By default, the numerically lowest local address is chosen. For example, if the addresses **172.16.1.1/12**, **172.16.1.2/12**, and **172.16.1.3/12** are configured on the same interface, the preferred address on the subnet (by default, **172.16.1.1**) would be used as a local address when you issue a **ping 172.16.1.5** command.

To set a different preferred address for the subnet, include the **preferred** statement:

preferred;

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family *family* address *address*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family* address *address*]

**Related
Documentation**

- [Configuring the Interface Address on page 36](#)
- *Junos OS Administration Library for Routing Devices*

Configuring Unicast RPF

For interfaces that carry IPv4 or IPv6 traffic, you can reduce the impact of denial of service (DoS) attacks by configuring unicast reverse path forwarding (RPF). Unicast RPF helps determine the source of attacks and rejects packets from unexpected source addresses on interfaces where unicast RPF is enabled.



NOTE: If you want to configure unicast RPF, your router must be equipped with the Internet Processor II application-specific integrated circuit (ASIC).

If you enable unicast RPF on live traffic, some packets are dropped while the packet forwarding components are updating.

For transit packets exiting the router through the tunnel, forwarding path features, such as RPF, forwarding table filtering, source class usage, and destination class usage are not supported on the interfaces you configure as the output interface for tunnel traffic. For firewall filtering, you must allow the output tunnel packets through the firewall filter applied to input traffic on the interface that is the next-hop interface towards the tunnel destination.

The following sections describe unicast RPF in detail:

- [Unicast RPF and Default Routes on page 53](#)
- [Unicast RPF with Routing Asymmetry on page 54](#)
- [Configuring Unicast RPF Strict Mode on page 54](#)
- [Configuring Unicast RPF Loose Mode on page 56](#)
- [Configuring Unicast RPF Loose Mode with Ability to Discard Packets on page 58](#)
- [Configuring Unicast RPF on a VPN on page 59](#)
- [Example: Configuring Unicast RPF on page 60](#)

Unicast RPF and Default Routes

When the active route cannot be chosen from the routes in a routing table, the router chooses a default route. A default route is equivalent to an IP address of 0.0.0.0/0. If you configure a default route, and you configure unicast RPF on an interface that the default route uses, unicast RPF behaves differently than it does otherwise. For information about configuring default routes, see the *Junos OS Routing Protocols Library for Routing Devices*.

To determine whether the default route uses an interface, enter the **show route** command:

```
user@host> show route address
```

address is the next-hop address of the configured default route. The default route uses the interfaces shown in the output of the **show route** command.

The following sections describe how unicast RPF behaves when a default route uses an interface and when a default route does not use an interface:

- [Unicast RPF Behavior with a Default Route on page 53](#)
- [Unicast RPF Behavior Without a Default Route on page 54](#)

Unicast RPF Behavior with a Default Route

On all routers except those with MPCs and the MX80 router, unicast RPF behaves as follows if you configure a default route that uses an interface configured with unicast RPF:

- Loose mode—All packets are automatically accepted. For this reason, we recommend that you not configure unicast RPF loose mode on interfaces that the default route uses.
- Strict mode—The packet is accepted when the source address of the packet matches any of the routes (either default or learned) that can be originated from the interface. Note that routes can have multiple destinations associated with them; therefore, if one of the destinations matches the incoming interface of the packet, the packet is accepted.

On all routers with MPCs and the MX80 router, unicast RPF behaves as follows if you configure a default route that uses an interface configured with unicast RPF:

- Loose mode—All packets except the packets whose source is learned from the default route are accepted. All packets whose source is learned from the default route are dropped at the Packet Forwarding Engine. The default route is treated as if the route does not exist.
- Strict mode—The packet is accepted when the source address of the packet matches any of the routes (either default or learned) that can be originated from the interface. Note that routes can have multiple destinations associated with them; therefore, if one of the destinations matches the incoming interface of the packet, the packet is accepted.

On all routers, the packet is not accepted when either of the following is true:

- The source address of the packet does not match a prefix in the routing table.
- The interface does not expect to receive a packet with this source address prefix.

Unicast RPF Behavior Without a Default Route


If you do not configure a default route, or if the default route does not use an interface configured with unicast RPF, unicast RPF behaves as described in “[Configuring Unicast RPF Strict Mode](#)” on page 54 and “[Configuring Unicast RPF Loose Mode](#)” on page 56. To summarize, unicast RPF without a default route behaves as follows:

- Strict mode—The packet is not accepted when either of the following is true:
 - The packet has a source address that does not match a prefix in the routing table.
 - The interface does not expect to receive a packet with this source address prefix.
- Loose mode—The packet is not accepted when the packet has a source address that does not match a prefix in the routing table.

Unicast RPF with Routing Asymmetry

In general, we recommend that you not enable unicast RPF on interfaces that are internal to the network because internal interfaces are likely to have *routing asymmetry*. Routing asymmetry means that a packet’s outgoing and return paths are different. Routers in the core of the network are more likely to have asymmetric reverse paths than routers at the customer or provider edge. [Figure 3 on page 54](#) shows unicast RPF in an environment with routing asymmetry.

Figure 3: Unicast RPF with Routing Asymmetry



In [Figure 3 on page 54](#), if you enable unicast RPF on interface `so-0/0/0`, traffic destined for Router A is not rejected. If you enable unicast RPF on interface `so-1/0/1`, traffic from Router A is rejected.

If you need to enable unicast RPF in an asymmetric routing environment, you can use fail filters to allow the router to accept incoming packets that are known to be arriving by specific paths. For an example of a fail filter that accepts packets with a specific source and destination address, see “[Example: Configuring Unicast RPF](#)” on page 60.

Configuring Unicast RPF Strict Mode

In strict mode, unicast RPF checks whether the incoming packet has a source address that matches a prefix in the routing table, and whether the interface expects to receive a packet with this source address prefix.

If the incoming packet fails the unicast RPF check, the packet is not accepted on the interface. When a packet is not accepted on an interface, unicast RPF counts the packet and sends it to an optional fail filter. If the fail filter is not configured, the default action is to silently discard the packet.

The optional fail filter allows you to apply a filter to packets that fail the unicast RPF check. You can define the fail filter to perform any filter operation, including accepting, rejecting, logging, sampling, or policing.

When unicast RPF is enabled on an interface, Bootstrap Protocol (BOOTP) packets and Dynamic Host Configuration Protocol (DHCP) packets are not accepted on the interface. To allow the interface to accept BOOTP packets and DHCP packets, you must apply a fail filter that accepts all packets with a source address of **0.0.0.0** and a destination address of **255.255.255.255**. For a configuration example, see [“Example: Configuring Unicast RPF” on page 60](#).

For more information about unicast RPF, see the *Junos OS Routing Protocols Library for Routing Devices*. For more information about defining fail filters, see the *Routing Policy Feature Guide for Routing Devices*.

To configure unicast RPF, include the **rpf-check** statement:

```
rpf-check <fail-filter filter-name>;
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces *interface-name* unit *logical-unit-number* family (inet | inet6)]**
- **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family (inet | inet6)]**

Using unicast RPF can have several consequences when implemented with traffic filters:

- RPF fail filters are evaluated after input filters and before output filters.
- If you configure a filter counter for packets dropped by an input filter, and you want to know the total number of packets dropped, you must also configure a filter counter for packets dropped by the RPF check.
- To count packets that fail the RPF check and are accepted by the RPF fail filter, you must configure a filter counter.
- If an input filter forwards packets anywhere other than the inet.0 or inet6.0 routing tables, the unicast RPF check is not performed.
- If an input filter forwards packets anywhere other than the routing instance the input interface is configured for, the unicast RPF check is not performed.

Example: Configuring Unicast RPF Strict Mode

Configure unicast RPF strict mode, and apply a fail filter that allows the interface to accept BOOTP packets and DHCP packets. The filter accepts all packets with a source address of **0.0.0.0** and a destination address of **255.255.255.255**.

To configure unicast RPF in strict mode:

1. Configure the fail filter:

```
[edit firewall]
filter rpf-special-case-dhcp-bootp {
  term allow-dhcp-bootp {
    from {
      source-address {
        0.0.0.0/32;
      }
      address {
        255.255.255.255/32;
      }
    }
    then {
      count rpf-dhcp-bootp-traffic;
      accept;
    }
  }
  term default {
    then {
      log;
      reject;
    }
  }
}
```

2. Configure unicast RPF on interfaces:

```
[edit]
interfaces {
  so-0/0/0 {
    unit 0 {
      family inet {
        rpf-check fail-filter rpf-special-case-dhcp-bootp;
      }
    }
  }
}
```

3. Commit the configuration.

```
[edit]
commit;
```

Configuring Unicast RPF Loose Mode

By default, unicast RPF uses strict mode. Unicast RPF loose mode is similar to unicast RPF strict mode and has the same configuration restrictions. The only check in loose mode is whether the packet has a source address with a corresponding prefix in the routing table; loose mode does not check whether the interface expects to receive a packet with a specific source address prefix. If a corresponding prefix is not found, unicast RPF loose mode does not accept the packet. As in strict mode, loose mode counts the failed packet and optionally forwards it to a fail filter, which either accepts, rejects, logs, samples, or polices the packet.

To configure unicast RPF loose mode, include the **mode**:

```
mode loose;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family (inet | inet6) rpf-check <fail-filter *filter-name*>]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family (inet | inet6) rpf-check <fail-filter *filter-name*>]

Example: Configuring Unicast RPF Loose Mode

In this example, no special configuration beyond device initialization is required.

Configure unicast RPF loose mode, and apply a fail filter that allows the interface to accept BOOTP packets and DHCP packets. The filter accepts all packets with a source address of 0.0.0.0 and a destination address of 255.255.255.255.

To configure unicast RPF in strict mode:

1. Configure the fail filter:

```
[edit firewall]
filter rpf-special-case-dhcp-bootp {
  term allow-dhcp-bootp {
    from {
      source-address {
        0.0.0.0/32;
      }
      address {
        255.255.255.255/32;
      }
    }
    then {
      count rpf-dhcp-bootp-traffic;
      accept;
    }
  }
  term default {
    then {
      log;
      reject;
    }
  }
}
```

2. Configure unicast RPF on interfaces:

```
[edit]
interfaces {
  so-0/0/0 {
    unit 0 {
      family inet {
        rpf-check fail-filter rpf-special-case-dhcp-bootp;
        mode loose;
      }
    }
  }
}
```

```
    }  
  }  
}
```

3. Commit the configuration.

```
[edit]  
commit;
```

Configuring Unicast RPF Loose Mode with Ability to Discard Packets

Starting with Junos OS Release 12.1, unicast RPF loose mode has the ability to discard packets with the source address pointing to the discard interface. This feature is supported on MX Series routers and on T Series routers with Type 1 FPCs, Type 2 FPCs, and Type 3 FPCs. Using unicast RPF loose mode, along with Remote Triggered Black Hole (RTBH) filtering, provides an efficient way to discard packets coming from known attack sources. BGP policies in edge routers ensure that packets with untrusted source addresses have their next hop set to a discard route. When a packet arrives at the router with an untrusted source address, unicast RPF performs a route lookup of the source address. Because the source address route points to a discard next hop, the packet is dropped and a counter is incremented. This feature is supported on both IPv4 (inet) and IPv6 (inet6) address families.

To configure unicast RPF loose mode with the ability to discard packets, include the **rpf-loose-mode-discard family inet** statement at the **[edit forwarding-options]** hierarchy level:

```
rpf-loose-mode-discard {  
  family {  
    inet;  
  }  
}
```

Example: Configuring Unicast RPF Loose Mode with Ability to Discard Packets

In this example, no special configuration beyond device initialization is required.

Configure unicast RPF strict mode, and apply a fail filter that allows the interface to accept BOOTP packets and DHCP packets. The filter accepts all packets with a source address of **0.0.0.0** and a destination address of **255.255.255.255**.

To configure unicast RPF in strict mode:

1. Configure the fail filter:

```
[edit firewall]  
filter rpf-special-case-dhcp-bootp {  
  term allow-dhcp-bootp {  
    from {  
      source-address {  
        0.0.0.0/32;  
      }  
      address {  
        255.255.255.255/32;  
      }  
    }  
  }  
}
```

```

    }
  }
  then {
    count rpf-dhcp-bootp-traffic;
    accept;
  }
}
term default {
  then {
    log;
    reject;
  }
}
}
}

```

2. Configure unicast RPF on interfaces:

```

[edit]
interfaces {
  so-0/0/0 {
    unit 0 {
      family inet {
        rpf-check fail-filter rpf-special-case-dhcp-bootp;
        mode loose;
      }
    }
  }
}

```

3. Configure the ability to discard packets.

```

[edit]
forwarding-options {
  rpf-loose-mode-discard {
    family {
      inet;
    }
  }
}

```

4. Commit the configuration.

```

[edit]
commit;

```

Configuring Unicast RPF on a VPN

You can configure unicast RPF on a VPN interface by enabling unicast RPF on the interface and including the **interface** statement at the **[edit routing-instances *routing-instance-name*]** hierarchy level.

You can configure unicast RPF only on the interfaces you specify in the routing instance. This means the following:

- For Layer 3 VPNs, unicast RPF is supported on the CE router interface.
- Unicast RPF is not supported on core-facing interfaces.

- For virtual-router routing instances, unicast RPF is supported on all interfaces you specify in the routing instance.
- If an input filter forwards packets anywhere other than the routing instance the input interface is configured for, the unicast RPF check is not performed.

For more information about VPNs and virtual-router routing instances, see the *Junos OS VPNs Library for Routing Devices*. For more information about FBF, see the *Junos OS Routing Protocols Library for Routing Devices*.

Example: Configuring Unicast RPF on a VPN

Configure unicast RPF on a Layer 3 VPN interface:

```
[edit interfaces]
so-0/0/0 {
  unit 0 {
    family inet {
      rpf-check;
    }
  }
}
[edit routing-instance]
VPN-A {
  interface so-0/0/0.0;
}
```

Example: Configuring Unicast RPF

Configure unicast RPF strict mode, and apply a fail filter that allows the interface to accept BOOTP packets and DHCP packets. The filter accepts all packets with a source address of 0.0.0.0 and a destination address of 255.255.255.255.

```
[edit firewall]
filter rpf-special-case-dhcp-bootp {
  term allow-dhcp-bootp {
    from {
      source-address {
        0.0.0.0/32;
      }
      address {
        255.255.255.255/32;
      }
    }
    then {
      count rpf-dhcp-bootp-traffic;
      accept;
    }
  }
  term default {
    then {
      log;
      reject;
    }
  }
}
```

```
[edit]
interfaces {
  so-0/0/0 {
    unit 0 {
      family inet {
        rpf-check fail-filter rpf-special-case-dhcp-bootp;
      }
    }
  }
}
```

- Related Documentation**
- [unicast-reverse-path](#)
 - [Example: Configuring Unicast Reverse-Path-Forwarding Check on page 64](#)

Configuring Targeted Broadcast

The following sections explain how to configure targeted broadcast on an egress interface and its options:

- [Configuring Targeted Broadcast and Its Options on page 61](#)
- [Display Targeted Broadcast Configuration Options on page 62](#)

Configuring Targeted Broadcast and Its Options

You can configure targeted broadcast on an egress interface with different options. You can either allow the IP packets destined for a Layer 3 broadcast address to be forwarded on the egress interface and to send a copy of the IP packets to the Routing Engine or you can allow the IP packets to be forwarded on the egress interface only. Note that the packets are broadcast only if the egress interface is a LAN interface.

To configure targeted broadcast and its options:

1. Configure the physical interface.

```
[edit]
user@host# set interfaces interface-name
```

2. Configure the logical unit number at the `[edit interfaces interface-name` hierarchy level.

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

3. Configure the protocol family as inet at the `[edit interfaces interface-name unit interface-unit-number` hierarchy level.

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

4. Configure targeted broadcast at the `[edit interfaces interface-name unit interface-unit-number family inet` hierarchy level

```
[edit interfaces interface-name unit interface--unit-number family inet]
user@host# set targeted-broadcast
```

5. Specify one of the following options as per requirement:

- To allow IP packets destined for a Layer 3 broadcast address to be forwarded on the egress interface and to send a copy of the IP packets to the Routing Engine.

```
[edit interfaces interface-name unit interface-unit-number family inet
targeted-broadcast]
user@host# set forward-and-send-to-re
```

- To allow IP packets to be forwarded on the egress interface only.

```
[edit interfaces interface-name unit interface-unit-number family inet
targeted-broadcast]
user@host# set forward-only
```



NOTE: Targeted broadcast does not work when the targeted broadcast option `forward-and-send-to-re` and the traffic sampling option `sampling` are configured on the same egress interface of an M320 router, a T640 router, or an MX960 router. To overcome this scenario, you must either disable one of these options or enable the `sampling` option with the targeted broadcast option `forward-only` on the egress interface. For information about traffic sampling, see *Configuring Traffic Sampling*.

Display Targeted Broadcast Configuration Options

The following topics display targeted broadcast configuration with its various options:

- [Forward IP Packets On the Egress Interface and To the Routing Engine on page 62](#)
- [Forward IP Packets On the Egress Interface Only on page 62](#)

Forward IP Packets On the Egress Interface and To the Routing Engine

Purpose Display the configuration when targeted broadcast is configured on the egress interface to forward the IP packets on the egress interface and to send a copy of the IP packets to the Routing Engine.

Action To display the configuration run the **show** command at the **[edit interfaces *interface-name* unit *interface-unit-number* family inet]** where the interface name is `ge-2/0/0`, the unit value is set to `0`, the protocol family is set to `inet`.

```
[edit interfaces interface-name unit interface-unit-number family inet]
user@host# show
targeted-broadcast {
  forward-and-send-to-re;
}
```

Forward IP Packets On the Egress Interface Only

Purpose Display the configuration when targeted broadcast is configured on the egress interface to forward the IP packets on the egress interface only.

Action To display the configuration run the **show** command at the **[edit interfaces *interface-name* unit *interface-unit-number* family inet]** where the interface name is ge-2/0/0, the unit value is set to 0, the protocol family is set to inet.

```
[edit interfaces interface-name unit interface-unit-number family inet]
user@host#show
targeted-broadcast {
    forward-only;
}
```

- Related Documentation**
- [targeted-broadcast on page 170](#)
 - [Understanding Targeted Broadcast on page 63](#)

Understanding Targeted Broadcast

Targeted broadcast is a process of flooding a target subnet with Layer 3 broadcast IP packets originating from a different subnet. The intent of targeted broadcast is to flood the target subnet with the broadcast packets on a LAN interface without broadcasting to the entire network. Targeted broadcast is configured with various options on the egress interface of the router or switch and the IP packets are broadcast only on the LAN (egress) interface. Targeted broadcast helps you implement remote administration tasks such as backups and wake-on LAN (WOL) on a LAN interface, and supports virtual routing and forwarding (VRF) instances.

Regular Layer 3 broadcast IP packets originating from a subnet are broadcast within the same subnet. When these IP packets reach a different subnet, they are forwarded to the Routing Engine (to be forwarded to other applications). Because of this, remote administration tasks such as backups cannot be performed on a particular subnet through another subnet. As a workaround you can enable targeted broadcast, to forward broadcast packets that originate from a different subnet.

Layer 3 broadcast IP packets have a destination IP address that is a valid broadcast address for the target subnet. These IP packets traverse the network in the same way as unicast IP packets until they reach the destination subnet. In the destination subnet, if the receiving router has targeted broadcast enabled on the egress interface, the IP packets are forwarded to an egress interface and the Routing Engine or to an egress interface only. The IP packets are then translated into broadcast IP packets which flood the target subnet only through the LAN interface (if there is no LAN interface, the packets are discarded), and all hosts on the target subnet receive the IP packets. If targeted broadcast is not enabled on the receiving router, the IP packets are treated as regular Layer 3 broadcast IP packets and are forwarded to the Routing Engine. If targeted broadcast is enabled without any options, the IP packets are discarded.

Targeted broadcast can be configured to forward the IP packets only to an egress interface, which is helpful when the router is flooded with packets to process, or to both an egress interface and the Routing Engine.



NOTE: Targeted broadcast does not work when the targeted broadcast option `forward-and-send-to-re` and the traffic sampling option `sampling` are configured on the same egress interface of an M320 router, a T640 router, or an MX960 router. To overcome this scenario, you must either disable one of these options or enable the sampling option with the targeted broadcast option `forward-only` on the egress interface. For information about traffic sampling, see *Configuring Traffic Sampling*.



NOTE: Any firewall filter that is configured on the Routing Engine loopback interface (lo0) cannot be applied to IP packets that are forwarded to the Routing Engine as a result of a targeted broadcast. This is because broadcast packets are forwarded as flood next hop and not as local next hop traffic, and you can only apply a firewall filter to local next hop routes for traffic directed towards the Routing Engine.

**Related
Documentation**

- [Configuring Targeted Broadcast on page 61](#)
- [targeted-broadcast on page 170](#)

Example: Configuring Unicast Reverse-Path-Forwarding Check

- [Understanding Unicast Reverse Path Forwarding on page 64](#)
- [Example: Configuring Unicast Reverse-Path-Forwarding Check on page 65](#)

Understanding Unicast Reverse Path Forwarding

IP spoofing can occur during a denial-of-service (DoS) attack. IP spoofing allows an intruder to pass IP packets to a destination as genuine traffic, when in fact the packets are not actually meant for the destination. This type of spoofing is harmful because it consumes the destination's resources.

A unicast reverse-path-forwarding (RPF) check is a tool to reduce forwarding of IP packets that might be spoofing an address. A unicast RPF check performs a route table lookup on an IP packet's source address, and checks the incoming interface. The router or switch determines whether the packet is arriving from a path that the sender would use to reach the destination. If the packet is from a valid path, the router or switch forwards the packet to the destination address. If it is not from a valid path, the router or switch discards the packet. Unicast RPF is supported for the IPv4 and IPv6 protocol families, as well as for the virtual private network (VPN) address family.



NOTE: Reverse path forwarding is not supported on the interfaces you configure as tunnel sources. This affects only the transit packets exiting the tunnel.

Example: Configuring Unicast Reverse-Path-Forwarding Check

Unicast reverse path forwarding (RPF) helps protect against DoS and DDoS attacks by verifying the unicast source address of each packet that arrives on an ingress interface where unicast RPF is enabled.

This example shows how to help defend ingress interfaces against denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks by configuring unicast RPF to filter incoming traffic.

- [Requirements on page 65](#)
- [Overview on page 65](#)
- [Configuration on page 66](#)
- [Verification on page 71](#)

Requirements

In this example, no special configuration beyond device initialization is required.

Overview

Large amounts of unauthorized traffic such as attempts to flood a network with fake (bogus) service requests in a DoS attack can consume network resources and deny service to legitimate users. One way to help prevent DoS and DDoS attacks is to verify that incoming traffic originates from legitimate network sources.

Unicast RPF helps ensure that a traffic source is legitimate (authorized) by comparing the source address of each packet that arrives on an interface to the forwarding table entry for its source address. If the device uses the same interface that the packet arrived on to reply to the packet's source, this verifies that the packet originated from an authorized source, and the device forwards the packet. If the device does not use the same interface that the packet arrived on to reply to the packet's source, the packet might have originated from an unauthorized source, and the device discards the packet.


In this example, Device B has unicast RPF configured. Device A is using OSPF to advertise a prefix for the link that connects to Device D. OSPF is enabled on the links between Device B and Device C and the links between Device A and Device C, but not on the links between Device A and Device B. Therefore, Device B learns about the route to Device D through Device C.

If ingress filtering is used in an environment where DHCP or BOOTP is used, it should be ensured that the packets with a source address of 0.0.0.0 and a destination address of 255.255.255.255 are allowed to reach the relay agent in routers when appropriate.

This example also includes a fail filter. When a packet fails the unicast RPF check, the fail filter is evaluated to determine if the packet should be accepted anyway. The fail filter in this example allows Device B's interfaces to accept Dynamic Host Configuration Protocol (DHCP) packets. The filter accepts all packets with a source address of 0.0.0.0 and a destination address of 255.255.255.255.

[Figure 4 on page 66](#) shows the sample network.

Figure 4: Unicast RPF Sample Topoolgy



Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device A

```

set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces fe-0/0/2 unit 5 family inet address 10.0.0.5/30
set interfaces fe-0/0/1 unit 17 family inet address 10.0.0.17/30
set interfaces fe-0/1/1 unit 25 family inet address 10.0.0.25/30
set interfaces fe-1/1/1 unit 29 family inet address 10.0.0.29/30
set protocols ospf export send-direct
set protocols ospf area 0.0.0.0 interface fe-0/1/1.25
set protocols ospf area 0.0.0.0 interface fe-1/1/1.29
set policy-options policy-statement send-direct from protocol direct
set policy-options policy-statement send-direct from route-filter 10.0.0.16/30 exact
set policy-options policy-statement send-direct then accept

```

Device B

```

set interfaces fe-1/2/0 unit 2 family inet rpf-check fail-filter rpf-special-case-dhcp
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/1/1 unit 6 family inet rpf-check fail-filter rpf-special-case-dhcp
set interfaces fe-1/1/1 unit 6 family inet address 10.0.0.6/30
set interfaces fe-0/1/1 unit 9 family inet rpf-check fail-filter rpf-special-case-dhcp
set interfaces fe-0/1/1 unit 9 family inet address 10.0.0.9/30
set interfaces fe-0/1/0 unit 13 family inet rpf-check fail-filter rpf-special-case-dhcp
set interfaces fe-0/1/0 unit 13 family inet address 10.0.0.13/30
set protocols ospf area 0.0.0.0 interface fe-0/1/1.9
set protocols ospf area 0.0.0.0 interface fe-0/1/0.13
set routing-options forwarding-table unicast-reverse-path active-paths
set firewall filter rpf-special-case-dhcp term allow-dhcp from source-address 0.0.0.0/32
set firewall filter rpf-special-case-dhcp term allow-dhcp from destination-address 255.255.255.255/32
set firewall filter rpf-special-case-dhcp term allow-dhcp then count rpf-dhcp-traffic
set firewall filter rpf-special-case-dhcp term allow-dhcp then accept
set firewall filter rpf-special-case-dhcp term default then log
set firewall filter rpf-special-case-dhcp term default then reject

```

Device C

```

set interfaces fe-1/2/0 unit 10 family inet address 10.0.0.10/30
set interfaces fe-0/0/2 unit 14 family inet address 10.0.0.14/30
set interfaces fe-1/0/2 unit 21 family inet address 10.0.0.21/30

```

```

set interfaces fe-1/2/2 unit 26 family inet address 10.0.0.26/30
set interfaces fe-1/2/1 unit 30 family inet address 10.0.0.30/30
set protocols ospf area 0.0.0.0 interface fe-1/2/0.10
set protocols ospf area 0.0.0.0 interface fe-0/0/2.14
set protocols ospf area 0.0.0.0 interface fe-1/2/2.26
set protocols ospf area 0.0.0.0 interface fe-1/2/1.30

```

Device D set interfaces fe-1/2/0 unit 18 family inet address 10.0.0.18/30

Device E set interfaces fe-1/2/0 unit 22 family inet address 10.0.0.22/30

Configuring Device A

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device A:

1. Configure the interfaces.

```

[edit interfaces]
user@A# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30

user@A# set fe-0/0/2 unit 5 family inet address 10.0.0.5/30

user@A# set fe-0/0/1 unit 17 family inet address 10.0.0.17/30

user@A# set fe-0/1/1 unit 25 family inet address 10.0.0.25/30

user@A# set fe-1/1/1 unit 29 family inet address 10.0.0.29/30

```
2. Configure OSPF.

```

[edit protocols ospf]
user@A# set export send-direct
user@A# set area 0.0.0.0 interface fe-0/1/1.25
user@A# set area 0.0.0.0 interface fe-1/1/1.29

```
3. Configure the routing policy.

```

[edit policy-options policy-statement send-direct]
user@A# set from protocol direct
user@A# set from route-filter 10.0.0.16/30 exact
user@A# set then accept

```
4. If you are done configuring Device A, commit the configuration.

```

[edit]
user@A# commit

```

Configuring Device B

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device B:

1. Configure the interfaces.

```
[edit interfaces]
user@B# set fe-1/2/0 unit 2 family inet address 10.0.0.2/30

user@B# set fe-1/1/1 unit 6 family inet address 10.0.0.6/30

user@B# set fe-0/1/1 unit 9 family inet address 10.0.0.9/30

user@B# set fe-0/1/0 unit 13 family inet address 10.0.0.13/30
```
2. Configure OSPF.

```
[edit protocols ospf area 0.0.0.0]
user@B# set interface fe-0/1/1.9
user@B# set interface fe-0/1/0.13
```
3. Configure unicast RPF, and apply the optional fail filter.

```
[edit interfaces]
user@B# set fe-1/2/0 unit 2 family inet rpf-check fail-filter rpf-special-case-dhcp

user@B# set fe-1/1/1 unit 6 family inet rpf-check fail-filter rpf-special-case-dhcp

user@B# set fe-0/1/1 unit 9 family inet rpf-check fail-filter rpf-special-case-dhcp

user@B# set fe-0/1/0 unit 13 family inet rpf-check fail-filter rpf-special-case-dhcp
```
4. (Optional) Configure the fail filter that gets evaluated if a packet fails the RPF check.

```
[edit firewall filter rpf-special-case-dhcp]
user@B# set term allow-dhcp from source-address 0.0.0.0/32
user@B# set term allow-dhcp from destination-address 255.255.255.255/32
user@B# set term allow-dhcp then count rpf-dhcp-traffic
user@B# set term allow-dhcp then accept
user@B# set term default then log
user@B# set term default then reject
```
5. (Optional) Configure only active paths to be considered in the RPF check.
This is the default behavior.

```
[edit routing-options forwarding-table]
user@B# set unicast-reverse-path active-paths
```
6. If you are done configuring Device B, commit the configuration.

```
[edit]
user@B# commit
```

Results

Confirm your configuration by issuing the **show firewall**, **show interfaces**, **show protocols**, **show routing-options**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

Device A user@A# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 10.0.0.1/30;
    }
  }
}
fe-0/0/2 {
  unit 5 {
    family inet {
      address 10.0.0.5/30;
    }
  }
}
fe-0/0/1 {
  unit 17 {
    family inet {
      address 10.0.0.17/30;
    }
  }
}
fe-0/1/1 {
  unit 25 {
    family inet {
      address 10.0.0.25/30;
    }
  }
}
fe-1/1/1 {
  unit 29 {
    family inet {
      address 10.0.0.29/30;
    }
  }
}

user@A# show protocols
ospf {
  export send-direct;
  area 0.0.0.0 {
    interface fe-0/1/1.25;
    interface fe-1/1/1.29;
  }
}

user@A# show policy-options
policy-statement send-direct {
  from {

```

```
        protocol direct;
        route-filter 10.0.0.16/30 exact;
    }
    then accept;
}

Device B user@B# show firewall
filter rpf-special-case-dhcp {
    term allow-dhcp {
        from {
            source-address {
                0.0.0.0/32;
            }
            destination-address {
                255.255.255.255/32;
            }
            destination-address {
                255.255.255.255/32;
            }
        }
        then {
            count rpf-dhcp-traffic;
            accept;
        }
    }
    term default {
        then {
            log;
            reject;
        }
    }
}

user@B# show interfaces
fe-1/2/0 {
    unit 2 {
        family inet {
            rpf-check fail-filter rpf-special-case-dhcp;
            address 10.0.0.2/30;
        }
    }
}
fe-1/1/1 {
    unit 6 {
        family inet {
            rpf-check fail-filter rpf-special-case-dhcp;
            address 10.0.0.6/30;
        }
    }
}
fe-0/1/1 {
    unit 9 {
        family inet {
            rpf-check fail-filter rpf-special-case-dhcp;
            address 10.0.0.9/30;
        }
    }
}
```



```
}
fe-0/1/0 {
  unit 13 {
    family inet {
      rpf-check fail-filter rpf-special-case-dhcp;
      address 10.0.0.13/30;
    }
  }
}
```

user@B# show protocols

```
ospf {
  area 0.0.0.0 {
    interface fe-0/1/1.9;
    interface fe-0/1/0.13;
  }
}
```

user@B# show routing-options

```
forwarding-table {
  unicast-reverse-path active-paths;
}
```

Enter the configurations on Device C, Device D, and Device E, as shown in [“CLI Quick Configuration” on page 66](#).

Verification

Confirm that the configuration is working properly.

- [Confirm That Unicast RPF Is Enabled on page 71](#)
- [Confirm That the Source Addresses Are Blocked on page 72](#)
- [Confirm That the Source Addresses Are Unblocked on page 72](#)

Confirm That Unicast RPF Is Enabled

Purpose Make sure that the interfaces on Device B have unicast RPF enabled.

Action user@B> show interfaces fe-0/1/0.13 extensive
Logical interface fe-0/1/0.13 (Index 73) (SNMP ifIndex 553) (Generation 208)
Flags: SNMP-Traps 0x4000 Encapsulation: ENET2
Traffic statistics:
Input bytes : 999390
Output bytes : 1230122
Input packets: 12563
Output packets: 12613
Local statistics:
Input bytes : 998994
Output bytes : 1230122
Input packets: 12563
Output packets: 12613
Transit statistics:
Input bytes : 396 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
Protocol inet, MTU: 1500, Generation: 289, Route table: 22
Flags: Sendbcst-pkt-to-re, uRPF
RPF Failures: Packets: 0, Bytes: 0
Addresses, Flags: Is-Preferred Is-Primary
Destination: 10.0.0.12/30, Local: 10.0.0.13, Broadcast: 10.0.0.15,
Generation: 241

Meaning The uRPF flag confirms that unicast RPF is enabled on this interface.

Confirm That the Source Addresses Are Blocked

Purpose Use the ping command to make sure that Device B blocks traffic from unexpected source addresses.

Action From Device A, ping Device B's interfaces, using 10.0.0.17 as the source address.

```
user@A> ping 10.0.0.6 source 10.0.0.17
PING 10.0.0.6 (10.0.0.6): 56 data bytes
^C
--- 10.0.0.6 ping statistics ---
3 packets transmitted, 0 packets received, 100% packet loss
```

Meaning As expected, the ping operation fails.

Confirm That the Source Addresses Are Unblocked

Purpose Use the ping command to make sure that Device B does not block traffic when the RPF check is deactivated.

Action 1. Deactivate the RPF check on one of the interfaces.
2. Rerun the ping operation.

```
user@B> deactivate interfaces fe-1/1/1.6 family inet rpf-check
user@A> ping 10.0.0.6 source 10.0.0.17
PING 10.0.0.2 (10.0.0.2): 56 data bytes
64 bytes from 10.0.0.2: icmp_seq=0 ttl=63 time=1.316 ms
```

```
64 bytes from 10.0.0.2: icmp_seq=1 ttl=63 time=1.263 ms
^C
--- 10.0.0.2 ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.263/1.289/1.316/0.027 ms
```

Meaning As expected, the ping operation succeeds.

Related Documentation

- *Example: Enabling Indirect Next Hops on the Packet Forwarding Engine*

CHAPTER 3

Network Interfaces Configuration Statements and Hierarchy

- [\[edit firewall\] Hierarchy Level on page 75](#)
- [\[edit interfaces\] Hierarchy Level on page 76](#)
- [\[edit logical-systems\] Hierarchy Level on page 92](#)
- [\[edit protocols pppoe\] Hierarchy Level on page 97](#)

[\[edit firewall\] Hierarchy Level](#)

The following CoS statements can be configured at the **[edit firewall]** hierarchy level. This is not a comprehensive list of statements available at the **[edit firewall]** hierarchy level.

```
[edit firewall]
  atm-policer policer-name {
    cdvt rate;
    logical-interface-policer;
    max-burst-size max-burst-size;
    peak-rate rate;
    policing-action (discard | discard-tag | count);
    sustained-rate rate;
  }
  family family-name {
    filter filter-name {
      term term-name {
        from {
          match-conditions;
        }
        then {
          dscp 0;
          forwarding-class class-name;
          loss-priority (high | low);
          three-color-policer {
            (single-rate | two-rate) policer-name;
          }
        }
      }
    }
  }
  simple-filter filter-name {
```




NOTE: The *accounting-profile* statement is an exception to this rule. The *accounting-profile* statement can be configured at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level, but it cannot be configured at the [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*] hierarchy level.

```

interfaces {
  traceoptions {
    file filename <files number> <match regular-expression> <size size> <world-readable |
      no-world-readable> ;
    flag flag <disable>;
  }
  interface-name {
    accounting-profile name;
    aggregated-ether-options {
      (flow-control | no-flow-control);
      lacp {
        (active | passive);
        link-protection {
          disable;
          (revertive | non-revertive);
          periodic interval;
          system-priority priority;
        }
      }
      link-protection;
      link-speed speed;
      (loopback | no-loopback);
      mc-ae {
        chassis-id chassis-id;
        mc-ae-id mc-ae-id;
        mode (active-active | active-standby);
        redundancy-group group-id;
        status-control (active | standby);
      }
      minimum-links number;
      source-address-filter {
        mac-address;
      }
      (source-filtering | no-source-filtering);
    }
    shared-scheduler;
    aggregated-sonet-options {
      link-speed speed | mixed;
      minimum-links number;
    }
    atm-options {
      cell-bundle-size cells;
      ilmi;
      linear-red-profiles profile-name {
        high-plp-max-threshold percent;
        low-plp-max-threshold percent;
        queue-depth cells high-plp-threshold percent low-plp-threshold percent;
      }
    }
  }
}

```

```
mpls {
  pop-all-labels {
    required-depth number;
  }
}
pic-type (atm1 | atm2);
plp-to-clp;
promiscuous-mode {
  vpi vpi-identifier;
}
scheduler-maps map-name {
  forwarding-class class-name {
    epd-threshold cells plp1 cells;
    linear-red-profile profile-name;
    priority (high | low);
    transmit-weight (cells number | percent number);
  }
  vc-cos-mode (alternate | strict);
}
use-null-cw;
vpi vpi-identifier {
  maximum-vcs maximum-vcs;
  oam-liveness {
    down-count cells;
    up-count cells;
  }
  oam-period (seconds | disable);
  shaping {
    (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate
    burst length);
    queue-length number;
  }
}
}
clocking clock-source;
damping {
  enable;
  half-life seconds;
  max-suppress seconds;
  reuse number;
  suppress number;
}
data-input (system | interface interface-name);
dce;
serial-options {
  clock-rate rate;
  clocking-mode (dce | internal | loop);
  control-polarity (negative | positive);
  cts-polarity (negative | positive);
  dcd-polarity (negative | positive);
  dce-options {
    control-signal (assert | de-assert | normal);
    cts (ignore | normal | require);
    dcd (ignore | normal | require);
    dsr (ignore | normal | require);
    dtr signal-handling-option;
  }
}
```



```

    ignore-all;
    indication (ignore | normal | require);
    rts (assert | de-assert | normal);
    tm (ignore | normal | require);
}
dsr-polarity (negative | positive);
dte-options {
    control-signal (assert | de-assert | normal);
    cts (ignore | normal | require);
    dcd (ignore | normal | require);
    dsr (ignore | normal | require);
    dtr signal-handling-option;
    ignore-all;
    indication (ignore | normal | require);
    rts (assert | de-assert | normal);
    tm (ignore | normal | require);
}
dtr-circuit (balanced | unbalanced);
dtr-polarity (negative | positive);
encoding (nrz | nrzi);
indication-polarity (negative | positive);
line-protocol protocol;
loopback mode;
rts-polarity (negative | positive);
tm-polarity (negative | positive);
transmit-clock invert;
}
description text;
alias alias-name;
dialer-options {
    pool pool-name <priority priority>;
}
disable;
ds0-options {
    bert-algorithm algorithm;
    bert-error-rate rate;
    bert-period seconds;
    byte-encoding (nx56 | nx64);
    fcs (16 | 32);
    idle-cycle-flag (flags | ones);
    invert-data;
    loopback payload;
    start-end-flag (filler | shared);
}
e1-options {
    bert-error-rate rate;
    bert-period seconds;
    fcs (16 | 32);
    framing (g704 | g704-no-crc4 | unframed);
    idle-cycle-flag (flags | ones);
    invert-data;
    loopback (local | remote);
    start-end-flag (filler | shared);
    timeslots time-slot-range;
}
e3-options {

```

```
atm-encapsulation (direct | plcp);
bert-algorithm algorithm;
bert-error-rate rate;
bert-period seconds;
framing feet;
compatibility-mode (digital-link | kentrox | larscom) <subrate value>;
fcs (16 | 32);
framing (g.751 | g.832);
idle-cycle-flag (filler | shared);
invert-data;
loopback (local | remote);
(payload-scrambler | no-payload-scrambler);
start-end-flag (filler | shared);
(unframed | no-unframed);
}
encapsulation type;
es-options {
    backup-interface es-fpc/pic/port;
}
fastether-options {
    802.3ad aex;
    (flow-control | no-flow-control);
    ignore-l3-incompletes;
    ingress-rate-limit rate;
    (loopback | no-loopback);
    mpls {
        pop-all-labels {
            required-depth number;
        }
    }
    source-address-filter {
        mac-address;
    }
    (source-filtering | no-source-filtering);
}
flexible-vlan-tagging;
gigether-options {
    802.3ad aex;
    (asynchronous-notification | no-asynchronous-notification);
    (auto-negotiation | no-auto-negotiation) remote-fault <local-interface-online |
        local-interface-offline>;
    auto-reconnect seconds;
    (flow-control | no-flow-control);
    ignore-l3-incompletes;
    (loopback | no-loopback);
    mpls {
        pop-all-labels {
            required-depth number;
        }
    }
    no-auto-mdix;
    source-address-filter {
        mac-address;
    }
    (source-filtering | no-source-filtering);
    ethernet-switch-profile {
```

```
(mac-learn-enable | no-mac-learn-enable);
tag-protocol-id [ tpids ];
ethernet-policer-profile {
    input-priority-map {
        ieee802.1p premium [ values ];
    }
    output-priority-map {
        classifier {
            premium {
                forwarding-class class-name {
                    loss-priority (high | low);
                }
            }
        }
    }
}
policer cos-policer-name {
    aggregate {
        bandwidth-limit bps;
        burst-size-limit bytes;
    }
    premium {
        bandwidth-limit bps;
        burst-size-limit bytes;
    }
}
}
}
(gratuitous-arp-reply | no-gratuitous-arp-reply);
hold-time up milliseconds down milliseconds;
ima-group-options {
    differential-delay number;
    frame-length (32 | 64 | 128 | 256);
    frame-synchronization {
        alpha number;
        beta number;
        gamma number;
    }
    minimum-links number;
    symmetry (symmetrical-config-and-operation |
        symmetrical-config-asymmetrical-operation);
    test-procedure {
        ima-test-start;
        ima-test-stop;
        interface name;
        pattern number;
        period number;
    }
    transmit-clock (common | independent);
    version (1.0 | 1.1);
}
ima-link-options group-id group-id;
interface-set interface-set-name {
    interface ethernet-interface-name {
        (unit unit-number | vlan-tags-outer vlan-tag);
    }
}
```

```
interface interface-name {
    (unit unit-number);
}
}
isdn-options {
    bchannel-allocation (ascending | descending);
    calling-number number;
    pool pool-name <priority priority>;
    spid1 spid-string;
    spid2 spid-string;
    static-tei-val value;
    switch-type (att5e | etsi | nil | ntdms100 | ntt);
    t310 seconds;
    tei-option (first-call | power-up);
}
keepalives <down-count number> <interval seconds> <up-count number>;
link-mode mode;
lmi {
    lmi-type (ansi | itu | c-lmi);
    n391dte number;
    n392dce number;
    n392dte number;
    n393dce number;
    n393dte number;
    t391dte seconds;
    t392dce seconds;
}
lsq-failure-options {
    no-termination-request;
    [ trigger-link-failure interface-name ];
}
mac mac-address;
mlfr-uni-nni-bundle-options {
    acknowledge-retries number;
    acknowledge-timer milliseconds;
    action-red-differential-delay (disable-tx | remove-link);
    drop-timeout milliseconds;
    fragment-threshold bytes;
    cisco-interoperability send-lip-remove-link-for-link-reject;
    hello-timer milliseconds;
    link-layer-overhead percent;
    lmi-type (ansi | itu | c-lmi);
    minimum-links number;
    mrru bytes;
    n391 number;
    n392 number;
    n393 number;
    red-differential-delay milliseconds;
    t391 seconds;
    t392 seconds;
    yellow-differential-delay milliseconds;
}
modem-options {
    dialin (console | routable);
    init-command-string initialization-command-string;
}
```

```

mtu bytes;
multi-chassis-protection {
  peer a.b.c.d {
    interface interface-name;
  }
}
multiservice-options {
  (core-dump | no-core-dump);
  (syslog | no-syslog);
}
native-vlan-id number;
no-gratuitous-arp-request;
no-keepalives;
no-partition {
  interface-type type;
}
no-vpivci-swapping;
otn-options {
  fec (efec | gfec | none);
  (laser-enable | no-laser-enable);
  (line-loopback | no-line-loopback);
  pass-thru;
  rate (fixed-stuff-bytes | no-fixed-stuff-bytes | pass-thru);
  transmit-payload-type number;
  trigger (oc-lof | oc-lom | oc-los | oc-wavelength-lock | odu-ais | odu-bbe-th | odu-bdi
    | odu-es-th | odu-lck | odu-oci | odu-sd | odu-ses-th | odu-ttim | odu-uas-th |
    opu-ptm | otu-ais | otu-bbe-th | otu-bdi | otu-es-th | otu-fec-deg | otu-fec-exe |
    otu-iae | otu-sd | otu-ses-th | otu-ttim | otu-uas-th);
  tti;
}
optics-options {
  wavelength nm;
  alarm alarm-name {
    (syslog | link-down);
  }
  warning warning-name {
    (syslog | link-down);
  }
}
partition partition-number oc-slice oc-slice-range interface-type type;
timeslots time-slot-range;
passive-monitor-mode;
per-unit-scheduler;
ppp-options {
  chap {
    access-profile name;
    default-chap-secret name;
    local-name name;
    passive;
  }
  compression {
    acfc;
    pfc;
  }
  dynamic-profile profile-name;
  no-termination-request;
}

```

```
pap {
    access-profile name;
    local-name name;
    local-password password;
    compression;
}
}
psn-vcip psn-vci-identifier;
psn-vpip psn-vpi-identifier;
receive-bucket {
    overflow (discard | tag);
    rate percentage;
    threshold bytes;
}
redundancy-options {
    priority sp-fpc/pic/port;
    secondary sp-fpc/pic/port;
    hot-standby;
}
satop-options {
    payload-size n;
}
schedulers number;
serial-options {
    clock-rate rate;
    clocking-mode (dce | internal | loop);
    control-polarity (negative | positive);
    cts-polarity (negative | positive);
    dcd-polarity (negative | positive);
    dce-options {
        control-signal (assert | de-assert | normal);
        cts (ignore | normal | require);
        dcd (ignore | normal | require);
        dsr (ignore | normal | require);
        dtr signal-handling-option;
        ignore-all;
        indication (ignore | normal | require);
        rts (assert | de-assert | normal);
        tm (ignore | normal | require);
    }
    dsr-polarity (negative | positive);
    dte-options {
        control-signal (assert | de-assert | normal);
        cts (ignore | normal | require);
        dcd (ignore | normal | require);
        dsr (ignore | normal | require);
        dtr signal-handling-option;
        ignore-all;
        indication (ignore | normal | require);
        rts (assert | de-assert | normal);
        tm (ignore | normal | require);
    }
}
dtr-circuit (balanced | unbalanced);
dtr-polarity (negative | positive);
encoding (nrz | nrzi);
indication-polarity (negative | positive);
```

```

line-protocol protocol;
loopback mode;
rts-polarity (negative | positive);
tm-polarity (negative | positive);
transmit-clock invert;
}
services-options {
  inactivity-timeout seconds;
  open-timeout seconds;
  session-limit {
    maximum number;
    rate new-sessions-per-second;
  }
  syslog {
    host hostname {
      facility-override facility-name;
      log-prefix prefix-number;
      services priority-level;
    }
  }
}
shdsl-options {
  annex (annex-a | annex-b);
  line-rate line-rate;
  loopback (local | remote);
  snr-margin {
    current margin;
    snext margin;
  }
}
sonet-options {
  aggregate asx;
  aps {
    advertise-interval milliseconds;
    annex-b;
    authentication-key key;
    fast-aps-switch;
    force;
    hold-time milliseconds;
    lockout;
    neighbor address;
    paired-group group-name;
    preserve-interface;
    protect-circuit group-name;
    request;
    revert-time seconds;
    switching-mode (bidirectional | unidirectional);
    working-circuit group-name;
  }
  bytes {
    c2 value;
    e1-quiet value;
    f1 value;
    f2 value;
    s1 value;
    z3 value;
  }
}

```

```
    z4 value;
  }
  fcs (16 | 32);
  loopback (local | remote);
  mpls {
    pop-all-labels {
      required-depth number;
    }
  }
  path-trace trace-string;
  (payload-scrambler | no-payload-scrambler);
  rfc-2615;
  trigger {
    defect ignore;
    hold-time up milliseconds down milliseconds;
  }
  vtmapping (itu-t | klm);
  (z0-increment | no-z0-increment);
}
speed (10m | 100m | 1g | oc3 | oc12 | oc48);
stacked-vlan-tagging;
switch-options {
  switch-port port-number {
    (auto-negotiation | no-auto-negotiation);
    speed (10m | 100m | 1g);
    link-mode (full-duplex | half-duplex);
  }
}
}
t1-options {
  bert-algorithm algorithm;
  bert-error-rate rate;
  bert-period seconds;
  buildout value;
  byte-encoding (nx56 | nx64);
  crc-major-alarm-threshold (1e-3 | 5e-4 | 1e-4 | 5e-5 | 1e-5);
  crc-minor-alarm-threshold (1e-3 | 5e-4 | 1e-4 | 5e-5 | 1e-5 | 5e-6 | 1e-6);
  fcs (16 | 32);
  framing (esf | sf);
  idle-cycle-flag (flags | ones);
  invert-data;
  line-encoding (ami | b8zs);
  loopback (local | payload | remote);
  remote-loopback-respond;
  start-end-flag (filler | shared);
  timeslots time-slot-range;
}
t3-options {
  atm-encapsulation (direct | plcp);
  bert-algorithm algorithm;
  bert-error-rate rate;
  bert-period seconds;
  buildout feet;
  (cbit-parity | no-cbit-parity);
  compatibility-mode (adtran | digital-link | kentrox | larscom | verilink) <subrate
    value>;
  fcs (16 | 32);
```



```

(feac-loop-respond | no-feac-loop-respond);
idle-cycle-flag value;
(long-buildout | no-long-buildout);
(loop-timing | no-loop-timing);
loopback (local | payload | remote);
(mac | no-mac);
(payload-scrambler | no-payload-scrambler);
start-end-flag (filler | shared);
}
traceoptions {
    flag flag <flag-modifier> <disable>;
}
transmit-bucket {
    overflow discard;
    rate percentage;
    threshold bytes;
}
(traps | no-traps);
unidirectional;
vlan-tagging;
vlan-vci-tagging;
unit logical-unit-number {
    accept-source-mac {
        mac-address mac-address {
            policer {
                input cos-policer-name;
                output cos-policer-name;
            }
        }
    }
}
accounting-profile name;
advisory-options {
    downstream-rate rate;
    upstream-rate rate;
}
allow-any-vci;
atm-scheduler-map (map-name | default);
backup-options {
    interface interface-name;
}
bandwidth rate;
cell-bundle-size cells;
clear-dont-fragment-bit;
compression {
    rtp {
        f-max-period number;
        maximum-contexts number <force>;
        queues [ queue-numbers ];
        port {
            minimum port-number;
            maximum port-number;
        }
    }
}
compression-device interface-name;
copy-tos-to-outer-ip-header;

```

```
demux-destination family;  
demux-source family;  
demux-options {  
    underlying-interface interface-name;  
}  
description text;  
alias alias-name;  
interface {  
    l2tp-interface-id name;  
    (dedicated | shared);  
}  
dialer-options {  
    activation-delay seconds;  
    callback;  
    callback-wait-period time;  
    deactivation-delay seconds;  
    dial-string [ dial-string-numbers ];  
    idle-timeout seconds;  
    incoming-map {  
        caller (caller-id | accept-all);  
        initial-route-check seconds;  
        load-interval seconds;  
        load-threshold percent;  
        pool pool-name;  
        redial-delay time;  
        watch-list {  
            [ routes ];  
        }  
    }  
}  
disable;  
disable-mlppp-inner-ppp-pfc;  
dlci dlci-identifier;  
drop-timeout milliseconds;  
dynamic-call-admission-control {  
    activation-priority priority;  
    bearer-bandwidth-limit kilobits-per-second;  
}  
encapsulation type;  
epd-threshold cells plp1 cells;  
fragment-threshold bytes;  
inner-vlan-id-range start start-id end end-id;  
input-vlan-map {  
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);  
    inner-tag-protocol-id tpid;  
    inner-vlan-id number;  
    tag-protocol-id tpid;  
    vlan-id number;  
}  
interleave-fragments;  
inverse-arp;  
layer2-policer {  
    input-policer policer-name;  
    input-three-color policer-name;  
    output-policer policer-name;  
    output-three-color policer-name;
```

```

}
link-layer-overhead percent;
minimum-links number;
mrru bytes;
multicast-dlci dlci-identifier;
multicast-vci vpi-identifier.vci-identifier;
multilink-max-classes number;
multipoint;
oam-liveness {
    down-count cells;
    up-count cells;
}
oam-period (seconds | disable);
output-vlan-map {
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    tag-protocol-id tpid;
    vlan-id number;
}
passive-monitor-mode;
peer-unit unit-number;
plp-to-clp;
point-to-point;
ppp-options {
    chap {
        access-profile name;
        default-chap-secret name;
        local-name name;
        passive;
    }
    compression {
        acfc;
        pfc;
        pap;
        default-pap-password password;
        local-name name;
        local-password password;
        passive;
    }
    dynamic-profile profile-name;
    lcp-max-conf-req number;
    lcp-restart-timer milliseconds;
    loopback-clear-timer seconds;
    ncp-max-conf-req number;
    ncp-restart-timer milliseconds;
}
pppoe-options {
    access-concentrator name;
    auto-reconnect seconds;
    (client | server);
    service-name name;
    underlying-interface interface-name;
}
proxy-arp;
service-domain (inside | outside);

```

```
shaping {
  (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate
   burst length);
  queue-length number;
}
short-sequence;
transmit-weight number;
(traps | no-traps);
trunk-bandwidth rate;
trunk-id number;
tunnel {
  backup-destination address;
  destination address;
  key number;
  routing-instance {
    destination routing-instance-name;
  }
  source source-address;
  ttl number;
}
vci vpi-identifier.vci-identifier;
vci-range start start-vci end end-vci;
vpi vpi-identifier;
vlan-id number;
vlan-id-list [vlan-id vlan-id-vlan-id];
vlan-id-range number-number;
vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
vlan-tags-outer tpid.vlan-id inner-list [vlan-id vlan-id-vlan-id];
family family {
  accounting {
    destination-class-usage;
    source-class-usage {
      direction;
    }
  }
  access-concentrator name;
  address address {
    destination address;
  }
  bundle ml-fpc/pic/port | ls-fpc/pic/port);
  direct-connect;
  duplicate-protection;
  dynamic-profile profile-name;
  filter {
    group filter-group-number;
    input filter-name;
    input-list {
      [ filter-names ];
      output filter-name;
    }
    output-list {
      [ filter-names ];
    }
  }
  ipsec-sa sa-name;
  keep-address-and-control;
```

```

max-sessions number;
max-sessions-vs-a-ignore;
mtu bytes;
multicast-only;
negotiate-address;
no-redirects;
policer {
    arp policer-template-name;
    disable-arp-policer;
    input policer-template-name;
    output policer-template-name;
}
primary;
proxy inet-address address;
receive-options-packets;
receive-ttl-exceeded;
remote (inet-address address | mac-address address);
rpf-check {
    fail-filter filter-name;
    mode loose;
}
sampling {
    direction;
}
service {
    input {
        service-set service-set-name <service-filter filter-name>;
        post-service-filter filter-name;
    }
    output {
        service-set service-set-names <service-filter filter-name>;
    }
}
service-name-table table-name;
short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
    maximum-seconds>;
targeted-broadcast {
    forward-and-send-to-re;
    forward-only;
}
(translate-discard-eligible | no-translate-discard-eligible);
(translate-fecn-and-becn | no-translate-fecn-and-becn);
translate-plp-control-word-de;
unnumbered-address interface-name <destination address destination-profile
    profile-name | preferred-source-address address>;
address address {
    arp ip-address (mac | multicast-mac) mac-address <publish>;
    broadcast address;
    destination address;
    destination-profile name;
    eui-64;
    multipoint-destination address (dlci dlci-identifier | vci vci-identifier);
    multipoint-destination address {
        epd-threshold cells plp1 cells;
        inverse-arp;
        oam-liveness {

```

```
        up-count cells;  
        down-count cells;  
    }  
    oam-period (seconds | disable);  
    shaping {  
        (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained  
            rate burst length);  
        queue-length number;  
    }  
    vci vpi-identifier.vci-identifier;  
}  
preferred;  
primary;  
(vrrp-group | vrrp-inet6-group) group-number {  
    (accept-data | no-accept-data);  
    advertise-interval seconds;  
    authentication-type authentication;  
    authentication-key key;  
    fast-interval milliseconds;  
    (preempt | no-preempt) {  
        hold-time seconds;  
    }  
    priority-number number;  
    track {  
        priority-cost seconds;  
        priority-hold-time interface-name {  
            bandwidth-threshold bits-per-second {  
                priority;  
            }  
            interface priority;  
        }  
        route ip-address/mask routing-instance instance-name priority-cost cost;  
    }  
    virtual-address [ addresses ];  
}  
}  
}  
}
```

- Related Documentation**
- [Junos OS Hierarchy and RFC Reference](#)
 - [Ethernet Interfaces](#)
 - [Junos OS Network Interfaces Library for Routing Devices](#)

[edit logical-systems] Hierarchy Level

The following lists the statements that can be configured at the [edit logical-systems] hierarchy level that are also documented in this manual. For more information about logical systems, see the *Logical Systems Feature Guide for Routing Devices*.

```
logical-systems logical-system-name {  
    interfaces interface-name {
```

```

unit logical-unit-number {
  accept-source-mac {
    mac-address mac-address {
      policer {
        input cos-policer-name;
        output cos-policer-name;
      }
    }
  }
  allow-any-vci;
  atm-scheduler-map (map-name | default);
  bandwidth rate;
  backup-options {
    interface interface-name;
  }
  cell-bundle-size cells;
  clear-dont-fragment-bit;
  compression {
    rtp {
      f-max-period number;
      port {
        minimum port-number;
        maximum port-number;
      }
      queues [ queue-numbers ];
    }
  }
  compression-device interface-name;
  description text;
  interface {
    l2tp-interface-id name;
    (dedicated | shared);
  }
  dialer-options {
    activation-delay seconds;
    deactivation-delay seconds;
    dial-string [ dial-string-numbers ];
    idle-timeout seconds;
    initial-route-check seconds;
    load-threshold number;
    pool pool;
    remote-name remote-callers;
    watch-list {
      [ routes ];
    }
  }
  disable;
  dlci dlci-identifier;
  drop-timeout milliseconds;
  dynamic-call-admission-control {
    activation-priority priority;
    bearer-bandwidth-limit kilobits-per-second;
  }
  encapsulation type;
  epd-threshold cells plp1 cells;
  fragment-threshold bytes;

```

```
input-vlan-map {
    inner-tag-protocol-id;
    inner-vlan-id;
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
    tag-protocol-id tpid;
    vlan-id number;
}
interleave-fragments;
inverse-arp;
layer2-policer {
    input-policer policer-name;
    input-three-color policer-name;
    output-policer policer-name;
    output-three-color policer-name;
}
link-layer-overhead percent;
minimum-links number;
mrru bytes;
multicast-dlci dlci-identifier;
multicast-vci vpi-identifier.vci-identifier;
multilink-max-classes number;
multipoint;
oam-liveness {
    up-count cells;
    down-count cells;
}
oam-period (seconds | disable);
output-vlan-map {
    inner-tag-protocol-id;
    inner-vlan-id;
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-swap);
    tag-protocol-id tpid;
    vlan-id number;
}
passive-monitor-mode;
peer-unit unit-number;
plp-to-clp;
point-to-point;
ppp-options {
    chap {
        access-profile name;
        default-chap-secret name;
        local-name name;
        passive;
    }
    compression {
        acfc;
        pfc;
    }
}
dynamic-profile profile-name;
pap {
    default-pap-password password;
    local-name name;
    local-password password;
    passive;
```



```

    }
}
proxy-arp;
service-domain (inside | outside);
shaping {
    (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate
    burst length);
    queue-length number;
}
short-sequence;
transmit-weight number;
(traps | no-traps);
trunk-bandwidth rate;
trunk-id number;
tunnel {
    backup-destination address;
    destination address;
    key number;
    routing-instance {
        destination routing-instance-name;
    }
    source source-address;
    ttl number;
}
vci vpi-identifier.vci-identifier;
vlan-id number;
vlan-id-list [vlan-id vlan-id–vlan-id]
vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
vlan-tags outer tpid.vlan-id inner-list [vlan-id vlan-id–vlan-id]
vpi vpi-identifier;
family family {
    accounting {
        destination-class-usage;
        source-class-usage {
            direction;
        }
    }
}
bundle interface-name;
filter {
    group filter-group-number;
    input filter-name;
    input-list {
        [ filter-names ];
    }
    output filter-name;
    output-list {
        [ filter-names ];
    }
}
ipsec-sa sa-name;
keep-address-and-control;
mtu bytes;
multicast-only;
no-redirects;
policer {
    arp policer-template-name;

```

```
    input policer-template-name;
    output policer-template-name;
}
primary;
proxy inet-address address;
receive-options-packets;
receive-ttl-exceeded;
remote (inet-address address | mac-address address);
rpf-check <fail-filter filter-name> {
    <mode loose>;
}
sampling {
    direction;
}
service {
    input {
        service-set service-set-name <service-filter filter-name>;
        post-service-filter filter-name;
    }
    output {
        service-set service-set-name <service-filter filter-name>;
    }
}
(translate-discard-eligible | no-translate-discard-eligible);
(translate-ecbn-and-becn | no-translate-ecbn-and-becn);
unnumbered-address interface-name destination address destination-profile
    profile-name;
address address {
    arp ip-address (mac | multicast-mac) mac-address <publish>;
    broadcast address;
    destination address;
    destination-profile name;
    eui-64;
    multipoint-destination address (dlci dlci-identifier | vci vci-identifier);
    multipoint-destination address {
        epd-threshold cells plp1 cells;
        inverse-arp;
        oam-liveness {
            up-count cells;
            down-count cells;
        }
        oam-period (seconds | disable);
        shaping {
            (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained
                rate burst length);
            queue-length number;
        }
        vci vpi-identifier.vci-identifier;
    }
    preferred;
    primary;
    (vrrp-group | vrrp-inet6-group) group-number {
        (accept-data | no-accept-data);
        advertise-interval seconds;
        authentication-type authentication;
        authentication-key key;
```

Related Documentation

- *Junos OS Hierarchy and RFC Reference*
- *Ethernet Interfaces*
- *Junos OS Network Interfaces Library for Routing Devices*

Copyright © 2014, Juniper Networks, Inc. 97

```
traceoptions {  
  file <filename> <files number> <match regular-expression > <size maximum-file-size>  
    <world-readable | no-world-readable>;  
  filter {  
    aci regular-expression;  
    ari regular-expression;  
    service-name regular-expression;  
    underlying-interface interface-name;  
  }  
  flag flag;  
  level (all | error | info | notice | verbose | warning);  
  no-remote-trace;  
}
```

CHAPTER 4

Statement Summary

- [address](#) on page 102
- [aggregate \(Hierarchical Policer\)](#) on page 104
- [accounting](#) on page 105
- [arp \(Interfaces\)](#) on page 106
- [bandwidth-limit \(Hierarchical Policer\)](#) on page 107
- [broadcast](#) on page 108
- [bundle](#) on page 109
- [burst-size-limit \(Hierarchical Policer\)](#) on page 110
- [cbr](#) on page 111
- [demux0 \(Dynamic Interface\)](#) on page 112
- [destination \(IPCP\)](#) on page 113
- [destination \(Tunnels\)](#) on page 114
- [destination-class-usage](#) on page 115
- [destination-profile](#) on page 115
- [dynamic-profiles](#) on page 116
- [epd-threshold \(Logical Interface\)](#) on page 123
- [eui-64](#) on page 124
- [family \(Dynamic Standard Interface\)](#) on page 125
- [family](#) on page 127
- [fast-aps-switch](#) on page 131
- [filter](#) on page 132
- [forward-and-send-to-re](#) on page 133
- [forward-only](#) on page 133
- [hierarchical-policer](#) on page 134
- [if-exceeding](#) on page 135
- [input](#) on page 135
- [input-list](#) on page 136
- [interface-mode](#) on page 137

- [interfaces on page 138](#)
- [interfaces \(Static and Dynamic Subscribers\) on page 139](#)
- [inverse-arp on page 143](#)
- [ipsec-sa on page 143](#)
- [keep-address-and-control on page 144](#)
- [logical-systems on page 144](#)
- [mode \(Dynamic Profiles\) on page 145](#)
- [mode \(Interfaces\) on page 145](#)
- [mtu on page 146](#)
- [multicast-only on page 148](#)
- [multipoint-destination on page 149](#)
- [negotiate-address on page 150](#)
- [no-redirects on page 150](#)
- [oam-liveness on page 151](#)
- [oam-period on page 152](#)
- [output on page 153](#)
- [output-list on page 153](#)
- [policer \(Interface\) on page 154](#)
- [post-service-filter on page 155](#)
- [preferred on page 155](#)
- [preferred-source-address on page 156](#)
- [premium \(Hierarchical Policer\) on page 157](#)
- [primary \(Address on Interface\) on page 158](#)
- [protocols on page 158](#)
- [proxy on page 159](#)
- [queue-length on page 160](#)
- [receive-options-packets on page 160](#)
- [receive-ttl-exceeded on page 161](#)
- [remote on page 161](#)
- [rpf-check \(Dynamic Profiles\) on page 162](#)
- [rpf-check \(interfaces\) on page 163](#)
- [rtvbr on page 164](#)
- [sampling \(Interfaces\) on page 165](#)
- [service \(Logical Interfaces\) on page 166](#)
- [service-filter \(Interfaces\) on page 167](#)
- [service-set on page 167](#)
- [shaping on page 168](#)

- [source-class-usage](#) on page 169
- [targeted-broadcast](#) on page 170
- [then](#) on page 171
- [translate-discard-eligible](#) on page 171
- [translate-fecn-and-becn](#) on page 172
- [unit \(Dynamic Profiles Standard Interface\)](#) on page 173
- [unit](#) on page 176
- [unnumbered-address \(Demux\)](#) on page 182
- [unnumbered-address \(Dynamic Profiles\)](#) on page 183
- [unnumbered-address \(Ethernet\)](#) on page 185
- [unnumbered-address \(PPP\)](#) on page 186
- [vbr](#) on page 187
- [vci](#) on page 188
- [vlan-id \(Logical Port in Bridge Domain\)](#) on page 189
- [vlan-id-list \(Interface in Bridge Domain\)](#) on page 189

address

```

Syntax  address address {
        arp ip-address (mac | multicast-mac) mac-address <publish>;
        broadcast address;
        destination address;
        destination-profile name;
        eui-64;
        master-only;
        multipoint-destination address dlcid dlcid-identifier;
        multipoint-destination address {
            epd-threshold cells;
            inverse-arp;
            oam-liveness {
                up-count cells;
                down-count cells;
            }
            oam-period (disable | seconds);
            shaping {
                (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate burst
                 length);
                queue-length number;
            }
            vci vpi-identifier.vci-identifier;
        }
        primary;
        preferred;
        (vrrp-group | vrrp-inet6-group) group-number {
            (accept-data | no-accept-data);
            advertise-interval seconds;
            authentication-type authentication;
            authentication-key key;
            fast-interval milliseconds;
            (preempt | no-preempt) {
                hold-time seconds;
            }
            priority-number number;
            track {
                priority-cost seconds;
                priority-hold-time interface-name {
                    interface priority;
                    bandwidth-threshold bits-per-second {
                        priority;
                    }
                }
            }
            route ip-address/mask routing-instance instance-name priority-cost cost;
        }
        virtual-address [ addresses ];
    }
}

```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family *family*],
 [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*
 family *family*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure the interface address.

Options *address*—Address of the interface.

- In Junos OS Release 13.3 and later, when you configure an IPv6 host address and an IPv6 subnet address on an interface, the commit operation fails.
- In releases earlier than Junos OS Release 13.3, when you use the same configuration on an interface, the commit operation succeeds, but only one of the IPv6 addresses that was entered is assigned to the interface. The other address is not applied.

The remaining statements are explained separately.



NOTE: The `edit logical-systems` hierarchy is not available on QFabric systems.

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

- Related Documentation**
- [Configuring the Protocol Family on page 33](#)
 - *Junos OS Administration Library for Routing Devices*
 - *family*
 - [negotiate-address on page 150](#)
 - [unnumbered-address \(Ethernet\) on page 185](#)


aggregate (Hierarchical Policier)

Syntax	<pre>aggregate { if-exceeding { bandwidth-limit <i>bandwidth</i>; burst-size-limit <i>burst</i>; } then { discard; } }</pre>
Hierarchy Level	[edit firewall hierarchical-policer]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	On M40e, M120, and M320 (with FFPC and SFPC) edge routers and T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs, T4000 routers with Type 5 FPC and Enhanced Scaling Type 4 FPC, configure an aggregate hierarchical policer.
Options	Options are described separately.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Applying Policers on page 11• <i>Junos OS Class of Service Library for Routing Devices</i>

accounting

Syntax	<pre> accounting { destination-class-usage; source-class-usage { direction; } } </pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Enable IP packet counters on an interface. The statements are explained separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Enabling Source Class and Destination Class Usage on page 25

arp (Interfaces)

Syntax	<code>arp <i>ip-address</i> (mac multicast-mac) <i>mac-address</i> publish;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inetaddress <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inetaddress <i>address</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces only, configure Address Resolution Protocol (ARP) table entries, mapping IP addresses to MAC addresses.
Options	<p><i>ip-address</i>—IP address to map to the MAC address. The IP address specified must be part of the subnet defined in the enclosing address statement.</p> <p>mac <i>mac-address</i>—MAC address to map to the IP address. Specify the MAC address as six hexadecimal bytes in one of the following formats: <i>nnnn.nnnn.nnnn</i> or <i>nn:nn:nn:nn:nn:nn</i>. For example, 0011.2233.4455 or 00:11:22:33:44:55.</p> <p>multicast-mac <i>mac-address</i>—Multicast MAC address to map to the IP address. Specify the multicast MAC address as six hexadecimal bytes in one of the following formats: <i>nnnn.nnnn.nnnn</i> or <i>nn:nn:nn:nn:nn:nn</i>. For example, 0011.2233.4455 or 00:11:22:33:44:55.</p> <p>publish—(Optional) Have the router or switch reply to ARP requests for the specified IP address. If you omit this option, the router or switch uses the entry to reach the destination but does not reply to ARP requests.</p>
<hr/>	
<div> NOTE: The edit logical-systems hierarchy is not available on QFabric systems.</div> <hr/>	
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Static ARP Table Entries</i>• <i>Configuring Static ARP Entries</i>

bandwidth-limit (Hierarchical Policer)

Syntax	<code>bandwidth-limit <i>bps</i>;</code>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> firewall hierarchical-policer aggregate if-exceeding], [edit dynamic-profiles <i>profile-name</i> firewall hierarchical-policer premium if-exceeding], [edit firewall hierarchical-policer aggregate if-exceeding], [edit firewall hierarchical-policer premium if-exceeding]
Release Information	Statement introduced in Junos OS Release 9.5. Support at the [edit dynamic-profiles ... if-exceeding] hierarchy level introduced in Junos OS Release 11.4.
Description	On M40e, M120, and M320 (with FFPC and SFPC) edge routers; on MPCs hosted on MX Series routers; on T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs; and on T4000 routers with Type 5 FPC and Enhanced Scaling Type 4 FPC, configure the maximum average bandwidth for premium or aggregate traffic in a hierarchical policer.
Options	<i>bps</i> —You can specify the number of bits per second either as a decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000). Range: 32,000 through 50,000,000,000 (32,000 through 100,000,000,000 on MX Series and T Series routers)
Required Privilege Level	firewall —To view this statement in the configuration. firewall-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Hierarchical Policer Configuration Overview</i> • <i>Policer Bandwidth and Burst-Size Limits</i> • <i>Policer Color-Marking and Actions</i> • <i>Single Token Bucket Algorithm</i> • <i>Determining Proper Burst Size for Traffic Policers</i> • <i>aggregate (Hierarchical Policer)</i> • burst-size-limit (Hierarchical Policer) on page 110 • premium (Hierarchical Policer) on page 157

broadcast


Syntax	<code>broadcast address;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Set the broadcast address on the network or subnet. On a subnet you cannot specify a host address of 0, nor can you specify a broadcast address.
Default	The default broadcast address has a host portion of all ones.
Options	address —Broadcast address. The address must have a host portion of either all ones or all zeros. You cannot specify the addresses 0.0.0.0 or 255.255.255.255 .



NOTE: The edit logical-systems hierarchy is not available on QFabric systems.

Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Interface Address on page 36

bundle

Syntax	<code>bundle (ml-<i>fpc/pic/port</i> ls-<i>fpc/pic/port</i>);</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Associate the multilink interface with the logical interface it is joining. You can include this statement for the mlfr-end-to-end and mlfr-uni-nni protocol families only.
<div>  NOTE: For M Series routers and T Series routers, the following caveats apply: <ul style="list-style-type: none"> • Maximum supported throughput on the bundle interfaces is 45 Mbps. • Bundling of the logical interfaces under a T3 physical interface into the same or different bundles is not supported. </div>	
Options	ml-<i>fpc/pic/port</i> —Name of the multilink interface you are linking. ls-<i>fpc/pic/port</i> —Name of the link services interface you are linking.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Junos OS Services Interfaces Library for Routing Devices</i>

burst-size-limit (Hierarchical Policer)

Syntax	<code>burst-size-limit bytes;</code>
Hierarchy Level	[edit dynamic-profiles profile-name firewall hierarchical-policer aggregate if-exceeding], [edit dynamic-profiles profile-name firewall hierarchical-policer premium if-exceeding], [edit firewall hierarchical-policer aggregate if-exceeding], [edit firewall hierarchical-policer premium if-exceeding]
Release Information	Statement introduced in Junos OS Release 9.5. Support at the [edit dynamic-profiles ... if exceeding] hierarchy level introduced in Junos OS Release 11.4.
Description	On M40e, M120, and M320 (with FFPC and SFPC) edge routers; on MPCs hosted on MX Series routers; on T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs; and on T4000 routers with Type 5 FPC and Enhanced Scaling Type 4 FPC, configure the burst-size limit for premium or aggregate traffic in a hierarchical policer.
Options	bytes —Burst-size limit in bytes. The minimum recommended value is the maximum transmission unit (MTU) of the IP packets being policed. You can specify the value either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000). Range: 1500 through 2,147,450,880 (1500 through 100,000,000,000 on MPCs hosted on MX Series routers)
Required Privilege Level	firewall —To view this statement in the configuration. firewall-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Hierarchical Policer Configuration Overview</i>• <i>Policer Bandwidth and Burst-Size Limits</i>• <i>Policer Color-Marking and Actions</i>• <i>Single Token Bucket Algorithm</i>• <i>Determining Proper Burst Size for Traffic Policers</i>• <i>Hierarchical Policers</i>• <i>aggregate (Hierarchical Policer)</i>• bandwidth-limit (Hierarchical Policer) on page 107• premium (Hierarchical Policer) on page 157

cbr

Syntax	<code>cbr rate;</code>
Hierarchy Level	[edit interfaces at- <i>fpc/pic/port</i> atm-options vpi <i>vpi-identifier</i> shaping], [edit interfaces at- <i>fpc/pic/port</i> unit <i>logical-unit-number</i> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i> shaping], [edit interfaces at- <i>fpc/pic/ port</i> unit <i>logical-unit-number</i> shaping], [edit logical-systems <i>logical-system-name</i> interfaces at- <i>fpc/pic/port</i> unit <i>logical-unit-number</i> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i> shaping], [edit logical-systems <i>logical-system-name</i> interfaces at- <i>fpc/pic/port</i> unit <i>logical-unit-number</i> shaping]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For ATM encapsulation only, define a constant bit rate bandwidth utilization in the traffic-shaping profile.
Default	Unspecified bit rate (UBR); that is, bandwidth utilization is unlimited.
Options	rate —Peak rate, in bits per second (bps) or cells per second (cps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000). You can also specify a value in cells per second by entering a decimal number followed by the abbreviation c ; values expressed in cells per second are converted to bits per second by means of the formula 1 cps = 384 bps. For ATM1 and ATM2 OC3 interfaces, the maximum available rate is 100 percent of <i>line-rate</i> , or 135,600,000 bps. For ATM1 OC12 interfaces, the maximum available rate is 50 percent of <i>line-rate</i> , or 271,263,396 bps. For ATM2 IQ interfaces, the maximum available rate is 542,526,792 bps.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Defining the ATM Traffic-Shaping Profile • rtvbr on page 164 • shaping on page 168 • vbr on page 187

demux0 (Dynamic Interface)

Syntax demux0 {
 unit *logical-unit-number* {
 demux-options {
 underlying-interface *interface-name*
 }
 }
 family *family* {
 access-concentrator *name*;
 address *address*;
 demux-source {
 source-prefix;
 }
 direct-connect;
 duplicate-protection;
 dynamic-profile *profile-name*;
 filter {
 input *filter-name*;
 output *filter-name*;
 }
 mac-validate (loose | strict):
 max-sessions *number*;
 max-sessions-vsa-ignore;
 rpf-check {
 fail-filter *filter-name*;
 mode loose;
 }
 service-name-table *table-name*
 short-cycle-protection <lockout-time-min *minimum-seconds* lockout-time-max
 maximum-seconds>;
 unnumbered-address *interface-name* <preferred-source-address *address*>;
 }
 filter {
 input *filter-name*;
 output *filter-name*;
 }
 vlan-id *number*;
 }
 }

Hierarchy Level [edit [dynamic-profiles](#) *profile-name* [interfaces](#)]

Release Information Statement introduced in Junos OS Release 9.3.

Description Configure the logical demultiplexing (demux) interface in a dynamic profile.

Logical IP demux interfaces do not support IPv4 and IPv6 dual stack.

The remaining statements are explained separately.

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

- Related Documentation**
- *Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles*
 - *Demultiplexing Interface Overview*

destination (IPCP)

- Syntax** `destination address destination-profile profile-name;`
- Hierarchy Level** [edit interfaces *interface-name* unit *logical-unit-number* family inet unnumbered-address *interface-name*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet unnumbered-address *interface-name*]
- Release Information** Statement introduced before Junos OS Release 7.4.
- Description** For unnumbered interfaces with PPP encapsulation, specify the IP address of the remote interface.
- Options** *address*—IP address of the remote interface.
- The remaining statement is explained separately.
- Required Privilege Level** interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.
- Related Documentation**
- [Configuring IPCP Options on page 42](#)
 - [address on page 102](#)
 - [negotiate-address on page 150](#)
 - *Junos OS Administration Library for Routing Devices*

destination (Tunnels)

Syntax	<code>destination address;</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <i>family</i> inet address <i>address</i>],</code> <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> <i>family</i> inet <i>unnumbered-address</i></code> <code> <i>interface-name</i>],</code> <code>[edit interfaces <i>interface-name</i> <i>unit</i> <i>logical-unit-number</i> tunnel],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i></code> <code> <i>family</i> inet address <i>address</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i></code> <code> <i>family</i> inet <i>unnumbered-address</i> <i>interface-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> <i>unit</i> <i>logical-unit-number</i></code> <code> tunnel]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX Series.
Description	For encrypted, PPP-encapsulated, and tunnel interfaces, specify the remote address of the connection.
Options	<i>address</i> —Address of the remote side of the connection.
Required Privilege Level	<i>interface</i> —To view this statement in the configuration. <i>interface-control</i> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Interface Address on page 36• <i>Configuring Generic Routing Encapsulation Tunneling (CLI Procedure)</i>• <i>Junos OS Services Interfaces Library for Routing Devices</i>• <i>point-to-point</i>

destination-class-usage

Syntax	<code>destination-class-usage;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet accounting], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet accounting]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Enable packet counters on an interface that count packets that arrive from specific customers and are destined for specific prefixes on the provider core router.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Enabling Source Class and Destination Class Usage on page 25 • accounting on page 105 • source-class-usage on page 169

destination-profile

Syntax	<code>destination-profile <i>name</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet unnumbered-address <i>interface-name</i> destination <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet unnumbered-address <i>interface-name</i> destination <i>address</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For interfaces with PPP encapsulation, assign PPP properties to the remote destination end. You define the profile at the [edit access group-profile <i>name</i> ppp] hierarchy level.
Options	<i>name</i> —Profile name defined at the [edit access group-profile <i>name</i> ppp] hierarchy level.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring IPCP Options on page 42 • destination (IPCP) on page 113 • Junos OS Administration Library for Routing Devices

dynamic-profiles

```

Syntax  dynamic-profiles {
        profile-name {
            class-of-service {
                interfaces {
                    interface-name ;
                }
                unit logical-unit-number {
                    classifiers {
                        type (classifier-name | default);
                    }
                    output-traffic-control-profile (profile-name | $junos-cos-traffic-control-profile);
                    rewrite-rules {
                        dscp (rewrite-name | default);
                        dscp-ipv6 (rewrite-name | default);
                        ieee-802.1 (rewrite-name | default) vlan-tag (outer | outer-and-inner);
                        inet-precedence (rewrite-name | default);
                    }
                }
            }
        }
        scheduler-maps {
            map-name {
                forwarding-class class-name scheduler scheduler-name;
            }
        }
        schedulers {
            (scheduler-name) {
                buffer-size (seconds | percent percentage | remainder | temporal microseconds);
                drop-profile-map loss-priority (any | low | medium-low | medium-high | high)
                    protocol (any | non-tcp | tcp) drop-profile profile-name;
                excess-priority (low | high | $junos-cos-scheduler-excess-priority);
                excess-rate (percent percentage | percent $junos-cos-scheduler-excess-rate);
                overhead-accounting (shaping-mode) <bytes (byte-value)>;
                priority priority-level;
                shaping-rate (rate | predefined-variable);
                transmit-rate (percent percentage | rate | remainder) <exact | rate-limit>;
            }
        }
        traffic-control-profiles profile-name {
            delay-buffer-rate (percent percentage | rate | $junos-cos-delay-buffer-rate);
            excess-rate (percent percentage | proportion value | percent $junos-cos-excess-rate);
            guaranteed-rate (percent percentage | rate | $junos-cos-guaranteed-rate);
            overhead-accounting (shaping-mode) <bytes (byte-value)>;
            scheduler-map map-name;
            shaping-rate (rate | predefined-variable);
        }
    }
    firewall {
        family family {
            fast-update-filter filter-name {
                interface-specific;
                match-order [match-order];
            }
        }
    }

```

```

term term-name {
  from {
    match-conditions;
  }
  then {
    action;
    action-modifiers;
  }
  only-at-create;
}
}
firewall {
  family family {
    fast-update-filter filter-name {
      interface-specific;
      match-order [match-order];
      term term-name {
        from {
          match-conditions;
        }
        then {
          action;
          action-modifiers;
        }
        only-at-create;
      }
    }
    filter filter-name {
      interface-specific;
      term term-name {
        from {
          match-conditions;
        }
        then {
          action;
          action-modifiers;
        }
      }
    }
  }
  policer policer-name {
    filter-specific;
    if-exceeding {
      (bandwidth-limit bps | bandwidth-percent percentage);
      burst-size-limit bytes;
    }
    logical-bandwidth-policer;
    logical-interface-policer;
    physical-interface-policer;
    then {
      policer-action;
    }
  }
}
hierarchical-policer policer-name {
  aggregate {
    if-exceeding {
      bandwidth-limit-limit bps;
      burst-size-limit bytes;
    }
    then {

```

```
        policer-action;
    }
}
premium {
    if-exceeding {
        bandwidth-limit bps;
        burst-size-limit bytes;
    }
    then {
        policer-action;
    }
}
}
three-color-policer policer-name {
    action {
        loss-priority high then discard;
    }
    logical-interface-policer;
    single-rate {
        (color-aware | color-blind);
        committed-burst-size bytes;
        committed-information-rate bps;
        excess-burst-size bytes;
    }
    two-rate {
        (color-aware | color-blind);
        committed-burst-size bytes;
        committed-information-rate bps;
        peak-burst-size bytes;
        peak-information-rate bps;
    }
}
}
}
policy-options {
    prefix-list name {
        ip-addresses;
    }
}
}
}
interfaces interface-name {
    interface-set interface-set-name {
        interface interface-name {
            unit logical unit number {
                advisory-options {
                    downstream-rate rate;
                    upstream-rate rate;
                }
            }
        }
    }
}
}
unit logical-unit-number {
    auto-configure {
        agent-circuit-identifier {
            dynamic-profile profile-name;
        }
    }
}
```



```

    }
}
encapsulation (atm-ccc-cell-relay | atm-ccc-vc-mux | atm-cisco-nlpid |
atm-tcc-vc-mux | atm-mlppp-llc | atm-nlpid | atm-ppp-llc | atm-ppp-vc-mux |
atm-snap | atm-tcc-snap | atm-vc-mux | ether-over-atm-llc |
ether-vpls-over-atm-llc | ether-vpls-over-fr | ether-vpls-over-ppp | ethernet |
frame-relay-ccc | frame-relay-ppp | frame-relay-tcc | frame-relay-ether-type |
frame-relay-ether-type-tcc | multilink-frame-relay-end-to-end | multilink-ppp |
ppp-over-ether | ppp-over-ether-over-atm-llc | vlan-bridge | vlan-ccc | vlan-vci-ccc
| vlan-tcc | vlan-vpls);
family family {
    address address;
    filter {
        adf {
            counter;
            input-precedence precedence;
            not-mandatory;
            output-precedence precedence;
            rule rule-value;
        }
        input filter-name (
            precedence precedence;
        )
        output filter-name {
            precedence precedence;
        }
    }
}
rpf-check {
    fail-filter filter-name;
    mode loose;
}
service {
    input {
        service-set service-set-name {
            service-filter filter-name;
        }
        post-service-filter filter-name;
    }
    input-vlan-map {
        inner-tag-protocol-id tpid;
        inner-vlan-id number;
        (push | swap);
        tag-protocol-id tpid;
        vlan-id number;
    }
    output {
        service-set service-set-name {
            service-filter filter-name;
        }
    }
    output-vlan-map {
        inner-tag-protocol-id tpid;
        inner-vlan-id number;
        (pop | swap);
        tag-protocol-id tpid;
        vlan-id number;
    }
}

```

```
    }
  }
  unnumbered-address interface-name <preferred-source-address address>;
}
ppp-options {
  chap;
  pap;
}
vlan-id number;
vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
}
}
interfaces {
  demux0 {...}
}
interfaces {
  pp0 {...}
}
protocols {
  igmp {
    interface interface-name {
      accounting;
      disable;
      group-policy;
      immediate-leave;
      no-accounting;
      promiscuous-mode;
      ssm-map ssm-map-name;
      static {
        group group {
          source source;
        }
      }
      version version;
    }
  }
  mld {
    interface interface-name {
      disable;
      (accounting | no-accounting);
      group-policy;
      immediate-leave;
      oif-map;
      passive;
      ssm-map ssm-map-name;
      static {
        group multicast-group-address {
          exclude;
          group-count number;
          group-increment increment;
          source ip-address {
            source-count number;
            source-increment increment;
          }
        }
      }
    }
  }
  version version;
```

```

    }
  }
  router-advertisement {
    interface interface-name {
      current-hop-limit number;
      default-lifetime seconds;
      (managed-configuration | no-managed-configuration);
      max-advertisement-interval seconds;
      min-advertisement-interval seconds;
      (other-stateful-configuration | no-other-stateful-configuration);
      prefix prefix;
      reachable-time milliseconds;
      retransmit-timer milliseconds;
    }
  }
}
routing-instances routing-instance-name {
  interface interface-name;
  routing-options {
    access {
      route prefix {
        next-hop next-hop;
        metric route-cost;
        preference route-distance;
        tag route-tag;
      }
    }
    access-internal {
      route subscriber-ip-address {
        qualified-next-hop underlying-interface {
          mac-address address;
        }
      }
    }
    multicast {
      interface interface-name {
        no-qos-adjust;
      }
    }
  }
}
rib routing-table-name {
  access {
    route prefix {
      next-hop next-hop;
      metric route-cost;
      preference route-distance;
      tag route-tag;
    }
  }
  access-internal {
    route subscriber-ip-address {
      qualified-next-hop underlying-interface {
        mac-address address;
      }
    }
  }
}

```

```
    }  
  }  
}  
routing-options {  
  access {  
    route prefix {  
      next-hop next-hop;  
      metric route-cost;  
      preference route-distance;  
      tag route-tag;  
    }  
  }  
  access-internal {  
    route subscriber-ip-address {  
      qualified-next-hop underlying-interface {  
        mac-address address;  
      }  
    }  
  }  
}  
multicast {  
  interface interface-name {  
    no-qos-adjust;  
  }  
}  
}  
variables {  
  variable-name {  
    default-value default-value;  
    equals expression;  
    mandatory;  
    uid;  
    uid-reference;  
  }  
}  
}  
}
```

Hierarchy Level [\[edit\]](#)

Release Information Statement introduced in Junos OS Release 9.2.
Support at the **filter**, **policer**, **hierarchical-policer**, **three-color-policer**, and **policy options** hierarchy levels introduced in Junos OS Release 11.4.

Description Create dynamic profiles for use with DHCP or PPP client access.

Options *profile-name*—Name of the dynamic profile; string of up to 80 alphanumeric characters.

The remaining statements are explained separately.

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

- Related Documentation**
- *Configuring a Basic Dynamic Profile*
 - *Configuring Dynamic VLANs Based on Agent Circuit Identifier Information*
 - *Dynamic Profiles Overview*

epd-threshold (Logical Interface)

Syntax	<code>epd-threshold <i>cells</i> plp1 <i>cells</i>;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	<p>For ATM2 IQ interfaces only, define the early packet discard (EPD) threshold on a VC. The EPD threshold is a limit on the number of transmit packets that can be queued. Packets that exceed the limit are discarded. For interfaces configured in trunk mode, you can also configure dual EPD thresholds depending on the packet loss priorities (PLPs).</p>
Default	<p>Approximately 1 percent of the available cell buffers. If shaping is enabled, the default EPD threshold is proportional to the shaping rate according to the following formula:</p> $\text{default epd-threshold} = \text{number of buffers} * \text{shaping rate} / \text{line rate}$ <p>The minimum EPD threshold value is 48 cells. If the default EPD threshold formula results in an EPD threshold of less than 48 cells, the result will be ignored, and the minimum value of 48 cells will be used.</p>
Options	<p>cells—Maximum number of cells.</p> <p>Range: For 1-port and 2-port OC12 interfaces, 48 through 425,984 cells</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the ATM2 IQ EPD Threshold</i> • <i>Configuring Two EPD Thresholds per Queue</i>

eui-64

Syntax	eui-64;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>number</i> family inet6 address <i>address</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.3 for EX Series switches. Statement introduced in Junos OS Release 12.2 for the QFX Series.
Description	For interfaces that carry IP version 6 (IPv6) traffic, automatically generate the host number portion of interface addresses.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Interface Address on page 36

family (Dynamic Standard Interface)

```
Syntax  family family {
    access-concentrator name;
    address address;
    direct-connect;
    duplicate-protection;
    dynamic-profile profile-name;
    filter {
        adf {
            counter;
            input-precedence precedence;
            not-mandatory;
            output-precedence precedence;
            rule rule-value;
        }
        input filter-name {
            precedence precedence;
        }
        output filter-name {
            precedence precedence;
        }
    }
    mac-validate (loose | strict);
    max-sessions number;
    max-sessions-vsa-ignore;
    rpf-check {
        fail-filter filter-name;
        mode loose;
    }
    service {
        input {
            service-set service-set-name {
                service-filter filter-name;
            }
            post-service-filter filter-name;
        }
        output {
            service-set service-set-name {
                service-filter filter-name;
            }
        }
    }
    service-name-table table-name
    short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
        maximum-seconds>;
    unnumbered-address interface-name <preferred-source-address address>;
}
```

Hierarchy Level [edit [dynamic-profiles](#) *profile-name* [interfaces](#) *interface-name* [unit](#) *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.2.
Option **pppoe** introduced in Junos OS Release 11.2.

Description Configure protocol family information for the logical interface.



NOTE: Not all subordinate stanzas are available to every protocol family.

Options *family*—Protocol family:

- **inet**—IP version 4 suite
- **inet6**—IP version 6 suite
- **pppoe**—(MX Series routers with MPCs only) Point-to-Point Protocol over Ethernet
- **vpls**—Virtual private LAN service

The remaining statements are explained separately.

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

Related Documentation

- *Example: Configuring Static Routing on Logical Systems*
- [Configuring the Protocol Family on page 33](#)

family

```

Syntax  family family {
        accounting {
            destination-class-usage;
            source-class-usage {
                (input | output | input output);
            }
        }
        access-concentrator name;
        address address {
            ... the address subhierarchy appears after the main [edit interfaces interface-name unit
                logical-unit-number family family-name] hierarchy ...
        }
        bundle interface-name;
        core-facing;
        demux-destination {
            destination-prefix;
        }
        demux-source {
            source-prefix;
        }
        direct-connect;
        duplicate-protection;
        dynamic-profile profile-name;
        filter {
            group filter-group-number;
            input filter-name;
            input-list [ filter-names ];
            output filter-name;
            output-list [ filter-names ];
        }
        interface-mode (access | trunk);
        ipsec-sa sa-name;
        keep-address-and-control;
        mac-validate (loose | strict);
        max-sessions number;
        max-sessions-vsa-ignore;
        mtu bytes;
        multicast-only;
        negotiate-address;
        no-redirects;
        policer {
            arp policer-template-name;
            input policer-template-name;
            output policer-template-name;
        }
        primary;
        protocols [inet iso mpls];
        proxy inet-address address;
        receive-options-packets;
        receive-ttl-exceeded;
        remote (inet-address address | mac-address address);
        rpf-check {

```

```
fail-filter filter-name
mode loose;
}
sampling {
input;
output;
}
service {
input {
post-service-filter filter-name;
service-set service-set-name <service-filter filter-name>;
}
output {
service-set service-set-name <service-filter filter-name>;
}
}
service-name-table table-name
short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
maximum-seconds>;
(translate-discard-eligible | no-translate-discard-eligible);
(translate-fecn-and-becn | no-translate-fecn-and-becn);
translate-plp-control-word-de;
unnumbered-address interface-name destination address destination-profile profile-name;
vlan-id number;
vlan-id-list [number number-number];
address address {
arp ip-address (mac | multicast-mac) mac-address <publish>;
broadcast address;
destination address;
destination-profile name;
eui-64;
master-only;
multipoint-destination address dlci dlci-identifier;
multipoint-destination address {
epd-threshold cells;
inverse-arp;
oam-liveness {
up-count cells;
down-count cells;
}
oam-period (disable | seconds);
shaping {
(cbr rate | rtvbr burst length peak rate sustained rate | vbr burst length peak rate
sustained rate);
queue-length number;
}
vci vpi-identifier.vci-identifier;
}
preferred;
primary;
vrrp-group group-id {
(accept-data | no-accept-data);
advertise-interval seconds;
authentication-key key;
authentication-type authentication;
fast-interval milliseconds;
```

```

    (preempt | no-preempt) {
        hold-time seconds;
    }
    priority number;
    track {
        interface interface-name {
            bandwidth-threshold bits-per-second priority-cost priority;
            priority-cost priority;
        }
        priority-hold-time seconds;
        route prefix routing-instance instance-name priority-cost priority;
    }
    }
    virtual-address [ addresses ];
    }
    virtual-link-local-address ipv6-address;
    }
}

```

Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Option max-sessions-vs-a-ignore introduced in Junos OS Release 11.4.
Description	Configure protocol family information for the logical interface.



NOTE: Not all subordinate stanzas are available to every protocol family.

fast-aps-switch

Syntax	fast-aps-switch;
Hierarchy Level	[edit interfaces <i>interface-name</i> sonet-options aps]
Release Information	Statement introduced in Junos OS Release 12.1.
Description	(M320 routers with Channelized OC3/STM1 Circuit Emulation PIC with SFP only and EX Series switches) Reduce the Automatic Protection Switching (APS) switchover time in Layer 2 circuits.




NOTE:

- Configuring this statement reduces the APS switchover time only when the Layer 2 circuit encapsulation type for the interface receiving traffic from a Layer 2 circuit neighbor is SAToP.
- When the fast-aps-switch statement is configured in revertive APS mode, you must configure an appropriate value for revert time to achieve reduction in APS switchover time.
- To prevent the logical interfaces in the data path from being shut down, configure appropriate hold-time values on all the interfaces in the data path that support TDM.
- The fast-aps-switch statement cannot be configured when the APS annex-b option is configured.
- The interfaces that have the fast-aps-switch statement configured cannot be used in virtual private LAN service (VPLS) environments.

Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Reducing APS Switchover Time in Layer 2 Circuits</i>

filter

Syntax	<pre>filter { group <i>filter-group-number</i>; input <i>filter-name</i>; input-list [<i>filter-names</i>]; output <i>filter-name</i>; output-list [<i>filter-names</i>]; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Apply a filter to an interface. You can also use filters for encrypted traffic. When you configure filters, you can configure them under the family ethernet-switching , inet , inet6 , mpls , or vpls only.
<div> NOTE: On QFX3500 and QFX3600 switches running Enhanced Layer 2 Software, VPLS is not supported.</div>	
Options	<p>group <i>filter-group-number</i>—Define an interface to be part of a filter group. Range: 1 through 255</p> <p>input <i>filter-name</i>—Name of one filter to evaluate when packets are received on the interface.</p> <p>output <i>filter-name</i>—Name of one filter to evaluate when packets are transmitted on the interface.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Applying a Filter to an Interface on page 20• <i>Junos OS Services Interfaces Library for Routing Devices</i>• <i>Routing Policy Feature Guide for Routing Devices</i>• <i>Junos OS Administration Library for Routing Devices</i>• <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i>• <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i>

- *Configuring Firewall Filters (CLI Procedure)*
- *Configuring Firewall Filters and Policers for VPLS*
- *family*
- *family*

forward-and-send-to-re

Syntax	forward-and-send-to-re;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet targeted-broadcast], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet targeted-broadcast]
Release Information	Statement introduced in Junos OS Release 10.2.
Description	Specify that IP packets destined for a Layer 3 broadcast address be forwarded to an egress interface and the Routing Engine. The packets are broadcast only if the egress interface is a LAN interface.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Targeted Broadcast on page 61 • targeted-broadcast on page 170 • Understanding Targeted Broadcast on page 63

forward-only

Syntax	forward-only;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet targeted-broadcast], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet targeted-broadcast]
Release Information	Statement introduced in Junos OS Release 10.2.
Description	Specify that IP packets destined for a Layer 3 broadcast address be forwarded to an egress interface only. The packets are broadcast only if the egress interface is a LAN interface.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Targeted Broadcast on page 61 • targeted-broadcast on page 170 • Understanding Targeted Broadcast on page 63

hierarchical-policer

Syntax hierarchical-policer *name* {
 aggregate {
 if-exceeding {
 bandwidth-limit *bandwidth*;
 burst-size-limit *burst*;
 }
 then {
 discard;
 }
 }
 premium {
 if-exceeding {
 bandwidth-limit *bandwidth*;
 burst-size-limit *burst*;
 }
 then {
 discard;
 }
 }
 }
 }

Hierarchy Level [edit firewall]

Release Information Statement introduced in Junos OS Release 9.5.

Description For M40e, M120, and M320 (with FFPC and SFPC) edge routers and T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs, specify a hierarchical policer.

Options Options are described separately.

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

Related Documentation

- [Applying Policers on page 11](#)
- *Junos OS Class of Service Library for Routing Devices*

if-exceeding

Syntax	if-exceeding { bandwidth-limit <i>bandwidth</i> ; burst-size-limit <i>burst</i> ; }
Hierarchy Level	[edit firewall hierarchical-policer aggregate], [edit firewall hierarchical-policer premium]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	For M40e, M120, and M320 (with FFPC and SFPC) edge routers and T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs, specify bandwidth and burst limits for an aggregate level of a hierarchical policer.
Options	Options are described separately.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Applying Policers on page 11 • <i>Junos OS Class of Service Library for Routing Devices</i>



input

Syntax	input { service-set <i>service-set-name</i> <service-filter <i>filter-name</i> >; post-service-filter <i>filter-name</i> ; }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet service], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet service]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Define one or more input service sets and filters, and one postservice filter to be applied to traffic.
Options	The remaining statements are explained separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Junos OS Services Interfaces Library for Routing Devices</i>

input-list

Syntax	<code>input-list [<i>filter-names</i>];</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> filter], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> filter]
Release Information	Statement introduced in Junos OS Release 7.6.
Description	Apply a group of filters to evaluate when packets are received on an interface.
Options	[<i>filter-names</i>] —Name of a filter to evaluate when packets are received on the interface. Up to 16 filters can be included in a filter input list.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Applying a Filter to an Interface on page 20• <i>Routing Policy Feature Guide for Routing Devices</i>• <i>Junos OS Administration Library for Routing Devices</i>• output-list on page 153

interface-mode

Syntax	interface-mode (access trunk);
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family ethernet-switching], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge]
Release Information	Statement introduced in Junos OS Release 9.2. Statement introduced in Junos OS Release 13.2X50-D10 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX Series.
Description	<div>  <p>NOTE: This statement supports the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that does not support ELS, see <i>port-mode</i>. For ELS details, see <i>Getting Started with Enhanced Layer 2 Software</i>.</p> </div> <p>(QFX Series 3500 and 3600 standalone switches)—Determine whether the logical interface accepts or discards packets based on VLAN tags. Specify the trunk option to accept packets with a VLAN ID that matches the list of VLAN IDs specified in the vlan-id or vlan-id-list statement, then forward the packet within the bridge domain or VLAN configured with the matching VLAN ID. Specify the access option to accept packets with no VLAN ID, then forward the packet within the bridge domain or VLAN configured with the VLAN ID that matches the VLAN ID specified in the vlan-id statement.</p> <div>  <p>NOTE: On MX Series routers, if you want IGMP snooping to be functional for a bridge domain, then you should not configure interface-mode and irb for that bridge. Such a configuration commit succeeds, but IGMP snooping is not functional, and a message informing the same is displayed. For more information, see <i>Configuring a Trunk Interface on a Bridge Network</i>.</p> </div>
Options	<p>access—Configure a logical interface to accept untagged packets. Specify the VLAN to which this interface belongs using the vlan-id statement.</p> <p>trunk—Configure a single logical interface to accept packets tagged with any VLAN ID specified with the vlan-id or vlan-id-list statement.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Configuring a Logical Interface for Access Mode Configuring a Logical Interface for Trunk Mode

- *Example: Connecting Access Switches to a Distribution Switch*

interfaces

Syntax	interfaces { ... }
Hierarchy Level	[edit]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure interfaces on the router or switch.
Default	The management and internal Ethernet interfaces are automatically configured. You must configure all other interfaces.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Physical Interface Configuration Statements Overview</i>• <i>Configuring Aggregated Ethernet Link Protection</i>

interfaces (Static and Dynamic Subscribers)

```

Syntax  interfaces {
        interface-name {
            unit logical-unit-number {
                auto-configure {
                    agent-circuit-identifier {
                        dynamic-profile profile-name;
                    }
                }
            }
            family family {
                access-concentrator name;
                address address;
                direct-connect;
                duplicate-protection;
                dynamic-profile profile-name;
                filter {
                    adf {
                        counter;
                        input-precedence precedence;
                        not-mandatory;
                        output-precedence precedence;
                        rule rule-value;
                    }
                    input filter-name (
                        precedence precedence;
                        shared-name filter-shared-name;
                    )
                    output filter-name {
                        precedence precedence;shared-name filter-shared-name;
                    }
                }
                max-sessions number;
                max-sessions-vs-a-ignore;
                rpf-check {
                    mode loose;
                }
                service {
                    input {
                        service-set service-set-name {
                            service-filter filter-name;
                        }
                        post-service-filter filter-name;
                    }
                    output {
                        service-set service-set-name {
                            service-filter filter-name;
                        }
                    }
                }
                service-name-table table-name
                short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
                    maximum-seconds>;
                unnumbered-address interface-name <preferred-source-address address>;
            }
        }
    }

```

```
    }
    filter {
        input filter-name;
        shared-name filter-shared-name;
        output filter-name;
        shared-name filter-shared-name;
    }
    ppp-options {
        chap;
        pap;
    }
    proxy-arp;
    vlan-id;
    vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
}
vlan-tagging;
}
interface-set interface-set-name {
    interface interface-name {
        unit logical unit number {
            advisory-options {
                downstream-rate rate;
                upstream-rate rate;
            }
        }
    }
}
pppoe-underlying-options {
    max-sessions number;
}
}
demux0 {
    unit logical-unit-number {
        demux-options {
            underlying-interface interface-name
        }
        family family {
            access-concentrator name;
            address address;
            direct-connect;
            duplicate-protection;
            dynamic-profile profile-name;
            demux-source {
                source-prefix;
            }
            filter {
                input filter-name {
                    precedence precedence;
                    shared-name filter-shared-name;
                }
                output filter-name {
                    precedence precedence;
                    shared-name filter-shared-name;
                }
            }
        }
        mac-validate (loose | strict);
        max-sessions number;
    }
}
```

```

max-sessions-vsa-ignore;
rpf-check {
    fail-filter filter-name;
    mode loose;
}
service-name-table table-name
short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
    maximum-seconds>;
unnumbered-address interface-name <preferred-source-address address>;
}
filter {
    input filter-name;
    output filter-name;
}
vlan-id number;
vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
}
}
pp0 {
    unit logical-unit-number {
        keepalives interval seconds;
        no-keepalives;
        pppoe-options {
            underlying-interface interface-name;
            server;
        }
        ppp-options {
            authentication [ authentication-protocols ];
            chap {
                challenge-length minimum minimum-length maximum maximum-length;
            }
            pap;
        }
        family inet {
            unnumbered-address interface-name;
            address address;
            service {
                input {
                    service-set service-set-name {
                        service-filter filter-name;
                    }
                    post-service-filter filter-name;
                }
                output {
                    service-set service-set-name {
                        service-filter filter-name;
                    }
                }
            }
        }
        filter {
            input filter-name {
                precedence precedence;
                shared-name filter-shared-name;
            }
            output filter-name {
                precedence precedence;
            }
        }
    }
}

```

```
        shared-name filter-shared-name;  
    }  
}  
}  
}  
}
```

Hierarchy Level [edit [dynamic-profiles](#) *profile-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Define interfaces for dynamic profiles.

Options *interface-name*—The interface variable (*\$junos-interface-ifd-name*). The interface variable is dynamically replaced with the interface the DHCP client accesses when connecting to the router.



NOTE: Though we do not recommend it, you can also enter the specific name of the interface you want to assign to the dynamic profile.

The remaining statements are explained separately.

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

Related Documentation

- *Configuring Static Subscriber Interfaces in Dynamic Profiles*
- *Configuring Dynamic Subscriber Interfaces Using IP Demux Interfaces in Dynamic Profiles*
- *Configuring Dynamic PPPoE Subscriber Interfaces Using Dynamic Profiles*
- *Configuring Dynamic VLANs Based on Agent Circuit Identifier Information*
- *Subscriber Interface Overview*
- *Relationship Between Subscribers and Interfaces in an Access Network*
- *Configuring Subscribers over Static Interfaces*
- *Demultiplexing Interface Overview*

inverse-arp

Syntax	<code>inverse-arp;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> multipoint-destination <i>destination</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> multipoint-destination <i>destination</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For ATM encapsulation, enable responses to receive inverse ATM ARP requests. For Frame Relay encapsulation, enable responses to receive inverse Frame Relay ARP requests.
Default	Inverse ARP is disabled on all ATM and Frame Relay interfaces.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Inverse ATM1 or ATM2 ARP</i> • <i>Configuring Inverse Frame Relay ARP</i>

ipsec-sa

Syntax	<code>ipsec-sa <i>sa-name</i>;</code>
Hierarchy Level	[edit interfaces <i>es-fpc/pic/port</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>es-fpc/pic/port</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the IP Security (IPsec) security association (SA) name associated with the interface.
Options	<i>sa-name</i> —IPsec security association name.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Junos OS Services Interfaces Library for Routing Devices</i> • <i>Junos OS Administration Library for Routing Devices</i>

keep-address-and-control

Syntax	keep-address-and-control;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>ccc</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>ccc</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For interfaces with encapsulation type PPP CCC, do not remove the address and control bytes before encapsulating the packet into a tunnel.
Default	If you do not include this statement, address and control bytes are removed before encapsulating the packet into a tunnel.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Disabling the Removal of Address and Control Bytes on page 10

logical-systems

Syntax	logical-systems { <i>logical-system-name</i> { ... <i>logical-system-configuration</i> ... } }
Hierarchy Level	[edit]
Release Information	Statement introduced before Junos OS Release 7.4. Statement name changed from logical-routers in Junos OS Release 9.3.
Description	Configure a logical system.
Options	<i>logical-system-name</i> —Name of the logical system.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Logical Systems Feature Guide for Routing Devices</i>

mode (Dynamic Profiles)

Syntax	mode loose;
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet) rpf-check]
Release Information	Statement introduced in Junos OS Release 9.6.
Description	Check whether the packet has a source address with a corresponding prefix in the routing table. If a corresponding prefix is not found, unicast reverse path forwarding (RPF) loose mode does not accept the packet. Unlike strict mode, loose mode does not check whether the interface expects to receive a packet with a specific source address prefix.
Default	If you do not include this statement, unicast RPF is in strict mode.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Unicast RPF Strict Mode on page 54

mode (Interfaces)

Syntax	mode loose;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) rpf-check], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) rpf-check]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Check whether the packet has a source address with a corresponding prefix in the routing table. If a corresponding prefix is not found, unicast reverse path forwarding (RPF) loose mode does not accept the packet. Unlike strict mode, loose mode does not check whether the interface expects to receive a packet with a specific source address prefix.
Default	If you do not include this statement, unicast RPF is in strict mode.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Unicast RPF Strict Mode on page 54

mtu

Syntax	<code>mtu bytes;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>],</p> <p>[edit interfaces <i>interface-range name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols l2circuit local-switching interface <i>interface-name</i> backup-neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i> backup-neighbor <i>address</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn interface <i>interface-name</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols vpls],</p> <p>[edit protocols l2circuit local-switching interface <i>interface-name</i> backup-neighbor <i>address</i>],</p> <p>[edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i>],</p> <p>[edit protocols l2circuit neighbor <i>address</i> interface <i>interface-name</i> backup-neighbor <i>address</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols l2vpn interface <i>interface-name</i>],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols vpls]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Support for Layer 2 VPNs and VPLS introduced in Junos OS Release 10.4.</p> <p>Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.</p> <p>Statement introduced in Junos OS Release 12.2 for ACX Series Universal Access Routers.</p> <p>Support at the [set interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>ccc</i>] hierarchy level introduced in Junos OS Release 12.3R3 for MX Series routers.</p>
Description	<p>Specify the maximum transmission unit (MTU) size for the media or protocol. The default MTU size depends on the device type. Changing the media MTU or protocol MTU causes an interface to be deleted and added again.</p> <p>To route jumbo data packets on an integrated routing and bridging (IRB) interface or routed VLAN interface (RVI) on EX Series switches, you must configure the jumbo MTU size on the member physical interfaces of the VLAN that you have associated with the IRB interface or RVI, as well as on the IRB interface or RVI itself (the interface named <i>irb</i> or <i>vlan</i>, respectively).</p>



CAUTION: For EX Series switches, setting or deleting the jumbo MTU size on an IRB interface or RVI while the switch is transmitting packets might cause packets to be dropped.



NOTE:

The MTU for an IRB interface is calculated by removing the Ethernet header overhead [6(DMAC)+6(SMAC)+2(EtherType)]. Because, the MTU is the lower value of the MTU configured on the IRB interface and the MTU configured on the IRB's associated bridge domain IFDs or IFLs, the IRB MTU is calculated as follows:

- In case of Layer 2 IFL configured with the `flexible-vlan-tagging` statement, the IRB MTU is calculated by including 8 bytes overhead (SVLAN+CVLAN).
- In case of Layer 2 IFL configured with the `vlan-tagging` statement, the IRB MTU is calculated by including a single VLAN 4 bytes overhead.



NOTE:

- If a packet whose size is larger than the configured MTU size is received on the receiving interface, the packet is eventually dropped. The value considered for MRU (maximum receive unit) size is also the same as the MTU size configured on that interface.
- Not all devices allow you to set an MTU value, and some devices have restrictions on the range of allowable MTU values. You cannot configure an MTU for management Ethernet interfaces (fxp0, em0, or me0) or for loopback, multilink, and multicast tunnel devices.
- On ACX Series routers, you can configure the protocol MTU by including the `mtu` statement at the [edit interfaces *interface-name* unit *logical-unit-number* family inet] or [edit interfaces *interface-name* unit *logical-unit-number* family inet6] hierarchy level.
 - If you configure the protocol MTU at any of these hierarchy levels, the configured value is applied to all families that are configured on the logical interface.
 - If you are configuring the protocol MTU for both inet and inet6 families on the same logical interface, you must configure the same value for both the families. It is not recommended to configure different MTU size values for inet and inet6 families that are configured on the same logical interface.

For more information about configuring MTU for specific interfaces and router or switch combinations, see *Configuring the Media MTU*.

Options	bytes —MTU size. Range: 256 through 9192 bytes, 256 through 9216 (EX Series switch interfaces), 256 through 9500 bytes (Junos OS 12.1X48R2 for PTX Series routers) Default: 1500 bytes (INET, INET6, and ISO families), 1448 bytes (MPLS), 1514 bytes (EX Series switch interfaces)
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i>• <i>Configuring Gigabit Ethernet Interfaces (CLI Procedure)</i>• <i>Configuring Routed VLAN Interfaces (CLI Procedure)</i>• <i>Configuring Integrated Routing and Bridging Interfaces (CLI Procedure)</i>• <i>Configuring the Media MTU</i>• <i>Configuring the MTU for Layer 2 Interfaces</i>• Setting the Protocol MTU on page 9

multicast-only

Syntax	multicast-only;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure the unit and family so that it can transmit and receive multicast traffic only. You can configure this property on the IP family only.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Protocol Family on page 33• <i>Junos OS Services Interfaces Library for Routing Devices</i>• <i>tunnel</i>

multipoint-destination

Syntax	<pre> multipoint-destination address dlcid dlcid-identifier; multipoint-destination address { epd-threshold cells; inverse-arp; oam-liveness { down-count cells; up-count cells; } oam-period (disable seconds); shaping { (cbr rate rtvbr peak rate sustained rate burst length vbr peak rate sustained rate burst length); queue-length number; } vci vpi-identifier.vci-identifier; } </pre>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>]</p>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For point-to-multipoint Frame Relay or ATM interfaces only, enable the support of multicast on the interface. You can configure multicast support on the interface if the Frame Relay or ATM switch performs multicast replication.
Options	<p>address—Address of the remote side of the point-to-multipoint connection.</p> <p>dlcid-identifier—For Frame Relay interfaces, the data-link connection identifier. Range: 0 through 0xFFFFFFF (24 bits)</p> <p>vci-identifier—For ATM interfaces, the virtual circuit identifier. Range: 0 through 16,384</p> <p>vpi-identifier—For ATM interfaces, the virtual path identifier. Range: 0 through 255 Default: 0</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Configuring a Point-to-Point ATM1 or ATM2 IQ Connection Configuring a Point-to-Multipoint Frame Relay Connection dlci encapsulation (Logical Interface)

negotiate-address

Syntax	negotiate-address;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For interfaces with PPP encapsulation, enable the interface to be assigned an IP address by the remote end.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring IPCP Options on page 42• address on page 102• unnumbered-address (PPP) on page 186• <i>Junos OS Administration Library for Routing Devices</i>

no-redirects

Syntax	no-redirects;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family family]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Do not send protocol redirect messages on the interface. To disable the sending of protocol redirect messages for the entire router or switch, include the no-redirects statement at the [edit system] hierarchy level.
Default	Interfaces send protocol redirect messages.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Disabling the Transmission of Redirect Messages on an Interface on page 10• <i>Junos OS Administration Library for Routing Devices</i>

oam-liveness

Syntax	<pre>oam-liveness { down-count <i>cells</i>; up-count <i>cells</i>; }</pre>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> atm-options vpi <i>vpi-identifier</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i> multipoint-destination <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i> multipoint-destination <i>address</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	<p>For ATM encapsulation only, configure Operation, Administration, and Maintenance (OAM) F5 loopback cell count thresholds. Not supported on ATM-over-SHDSL interfaces.</p> <p>For ATM2 IQ PICs only, configure OAM F4 loopback cell count thresholds at the [edit interfaces <i>interface-name</i> atm-options vpi <i>vpi-identifier</i>] hierarchy level.</p>
Options	<p>down-count <i>cells</i>—Minimum number of consecutive OAM F4 or F5 loopback cells lost before a VC is declared down. Range: 1 through 255 Default: 5 cells</p> <p>up-count <i>cells</i>—Minimum number of consecutive OAM F4 or F5 loopback cells received before a VC is declared up. Range: 1 through 255 Default: 5 cells</p>
Required Privilege Level	<p>interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Configuring the ATM OAM F5 Loopback Cell Threshold

oam-period

Syntax	oam-period (disable seconds);
Hierarchy Level	[edit interfaces <i>interface-name</i> atm-options vpi <i>vpi-identifier</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i> multipoint-destination <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i> multipoint-destination <i>address</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For ATM encapsulation only, configure the OAM F5 loopback cell period. Not supported on ATM-over-SHDLSL interfaces. For ATM2 IQ PICs only, configure the OAM F4 loopback cell period at the [edit interfaces <i>interface-name</i> atm-options vpi <i>vpi-identifier</i>] hierarchy level.
Default	If you omit this statement, OAM F5 loopback cells are not initiated, but the interface still responds if it receives OAM F5 loopback cells.
Options	disable —Disable the OAM loopback cell transmit feature. seconds —OAM loopback cell period. Range: 1 through 900 seconds
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">Defining the ATM OAM F5 Loopback Cell Period

output

Syntax	<code>output { service-set service-set-name <service-filter filter-name>; }</code>
Hierarchy Level	[edit interfaces interface-name unit logical-unit-number family inet service], [edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet service]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Define one or more output service sets and filters to be applied to traffic.
Options	The remaining statements are explained separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Junos OS Services Interfaces Library for Routing Devices</i>

output-list

Syntax	<code>output-list [filter-names];</code>
Hierarchy Level	[edit interfaces interface-name unit logical-unit-number family family filter], [edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family family filter]
Release Information	Statement introduced in Junos OS Release 7.6.
Description	Apply a group of filters to evaluate when packets are transmitted on an interface.
Options	[<i>filter-names</i>]—Name of a filter to evaluate when packets are transmitted on the interface. Up to 16 filters can be included in a filter input list.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Applying a Filter to an Interface on page 20 • input-list on page 136 • <i>Routing Policy Feature Guide for Routing Devices</i> • <i>Junos OS Services Interfaces Library for Routing Devices</i> • <i>Junos OS Administration Library for Routing Devices</i>


policer (Interface)

Syntax	<pre>policer { arp <i>policer-template-name</i>; disable-arp-policer input <i>policer-template-name</i>; output <i>policer-template-name</i>; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced before Junos OS Release 7.4. disable-arp-policer option added in Junos OS Release 14.1 for MX Series routers with MPCs (Trio-based FPCs).
Description	Apply a policer to an interface. To prevent the processing of ARP policers on the arriving ARP packets, include the disable-arp-policer statement. You can configure this statement only for interfaces with inet address families and on MX Series routers with MPCs. When you disable ARP policers per interface, the packets are continued to be policed by the distributed DoS (DDoS) ARP policer. The maximum rate of is 10000 pps per FPC.
Options	<p>arp <i>policer-template-name</i>—For inet family only, name of one policer to evaluate when ARP packets are received on the interface.</p> <p>disable-arp-policer—Configure the router to disable the processing of the specified ARP policers on the received ARP packets. Disabling ARP policers can cause denial-of-service (DoS) attacks on the system. Due to this possibility, we recommend that you exercise caution while disabling ARP policers.</p> <p>input <i>policer-template-name</i>—Name of one policer to evaluate when packets are received on the interface.</p> <p>output <i>policer-template-name</i>—Name of one policer to evaluate when packets are transmitted on the interface.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Applying Policers on page 11• <i>Configuring Firewall Filters and Policers for VPLS</i>• <i>Routing Policy Feature Guide for Routing Devices</i>• <i>Junos OS Services Interfaces Library for Routing Devices</i>

post-service-filter

Syntax	<code>post-service-filter <i>filter-name</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet service input], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet service input]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Define the filter to be applied to traffic after service processing. The filter is applied only if a service set is configured and selected.
Options	<i>filter-name</i> —Identifier for postservice filter.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Junos OS Services Interfaces Library for Routing Devices

preferred

Syntax	<code>preferred;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure this address to be the preferred address on the interface. If you configure more than one address on the same subnet, the preferred source address is chosen by default as the source address when you initiate frame transfers to destinations on the subnet.
<div style="display: flex; align-items: center;">  <div> <p>NOTE: The edit logical-systems hierarchy is not available on QFabric systems.</p> </div> </div>	
Default	The lowest-numbered address on the subnet is the preferred address.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Interface Address on page 36


preferred-source-address

Syntax	<code>preferred-source-address address;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> unnumbered-address <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> unnumbered-address <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.0.
Description	<p>For unnumbered Ethernet interfaces configured with a loopback interface as the donor interface, specify one of the loopback interface's secondary addresses as the preferred source address for the unnumbered Ethernet interface. Configuring the preferred source address enables you to use an IP address other than the primary IP address on some of the unnumbered Ethernet interfaces in your network.</p> <p>Configuration of a preferred source address for unnumbered Ethernet interfaces is supported for the IPv4 and IPv6 address families.</p>
Options	address —Secondary IP address of the donor loopback interface.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Preferred Source Address for Unnumbered Ethernet or Demux Interfaces on page 46• address on page 102• <i>Junos OS Administration Library for Routing Devices</i>

premium (Hierarchical Policer)

Syntax	<pre> premium { if-exceeding { bandwidth-limit <i>bandwidth</i>; burst-size-limit <i>burst</i>; } then { discard; } } </pre>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> firewall hierarchical-policer], [edit firewall hierarchical-policer]
Release Information	Statement introduced in Junos OS Release 9.5. Support at the [edit dynamic-profiles ... hierarchical-policer <i>name</i>] hierarchy level introduced in Junos OS Release 11.4.
Description	On M40e, M120, and M320 edge routers with FPC input as FFPC and FPC output as SFPC, and on MX Series, T320, T640, and T1600 edge routers with Enhanced Intelligent Queuing (IQE) PICs, T4000 routers with Type 5 FPC and Enhanced Scaling Type 4 FPC, specify a premium level for a hierarchical policer.
Options	Options are described separately.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Applying Policers on page 11 • <i>Guidelines for Applying Traffic Policers</i> • <i>Hierarchical Policer Configuration Overview</i> • <i>Hierarchical Policers</i> • <i>aggregate (Hierarchical Policer)</i> • bandwidth-limit (Hierarchical Policer) on page 107 • burst-size-limit (Hierarchical Policer) on page 110 • <i>hierarchical-policer</i> • <i>if-exceeding (Hierarchical Policer)</i>

primary (Address on Interface)

Syntax	<code>primary;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure this address to be the primary address of the protocol on the interface. If the logical unit has more than one address, the primary address is used by default as the source address when packet transfer originates from the interface and the destination address does not indicate the subnet.
<div> NOTE: The edit logical-systems hierarchy is not available on QFabric systems.</div>	
Default	For unicast traffic, the primary address is the lowest non-127 (in other words, non-loopback) preferred address on the unit.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Interface Address on page 36

protocols

Syntax	<code>protocols [inet iso mpls];</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>tcc</i>]
Release Information	Statement introduced in Junos OS Release 8.3.
Description	For Layer 2.5 VPNs on T Series, MX Series, M120, and M320 routers support, configure IS-IS (ISO traffic) or MPLS traffic to traverse a TCC interface. By default, IPv4 (inet) traffic runs on T Series, MX, Series, M120, and M320 routers and over TCC interfaces. You must configure the same traffic type on both ends of the Layer 2.5 VPN.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring IS-IS or MPLS Traffic for TCC Interfaces

proxy

Syntax	<code>proxy inet-address <i>address</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family tcc], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family tcc]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For Layer 2.5 VPNs using an Ethernet interface as the TCC router, configure the IP address for which the TCC router is proxying. Ethernet TCC is supported on interfaces that carry IPv4 traffic only. Ethernet TCC encapsulation is supported on 1-port Gigabit Ethernet, 2-port Gigabit Ethernet, 4-port Gigabit Ethernet, and 4-port Fast Ethernet PICs only. Ethernet TCC is not supported on the T640 router.
Options	inet-address —Configure the IP address of the neighbor to the TCC router.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Ethernet TCC</i> • <i>Example: Configuring an Ethernet TCC or Extended VLAN TCC</i> • remote on page 161 • <i>Junos OS VPNs Library for Routing Devices</i>

queue-length

Syntax	<code>queue-length <i>number</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i> shaping], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> shaping], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i> shaping], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> shaping]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	For ATM1 interfaces only, define the maximum queue length in the traffic-shaping profile. For ATM1 PICs, each VC has its own independent shaping parameters.
Default	Buffer usage is unregulated.
Options	<i>number</i> —Maximum number of packets the queue can contain. Range: 1 through 16,383 packets Default: 16,383 packets
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the ATM1 Queue Length</i>

receive-options-packets

Syntax	<code>receive-options-packets;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For a Monitoring Services PIC and an ATM or SONET/SDH PIC installed in an M160, M40e, or T Series router, guarantee conformity with cflowd records structure. This statement is required when you enable passive monitoring.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Enabling Passive Monitoring on ATM Interfaces</i>• <i>Enabling Passive Monitoring on SONET/SDH Interfaces</i>

receive-ttl-exceeded

Syntax	receive-ttl-exceeded;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For Monitoring Services PIC and an ATM or SONET/SDH PIC installed in an M160, M40e, or T Series router, guarantee conformity with cflowd records structure. This statement is required when you enable passive monitoring.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Enabling Passive Monitoring on ATM Interfaces</i> • <i>Enabling Passive Monitoring on SONET/SDH Interfaces</i>

remote

Syntax	remote { (inet-address <i>address</i> mac-address <i>address</i>); }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family tcc], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family tcc]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For Layer 2.5 VPNs using an Ethernet interface as the TCC router, configure the location of the remote router. Ethernet TCC is supported on interfaces that carry IPv4 traffic only. Ethernet TCC encapsulation is supported on 1-port Gigabit Ethernet, 2-port Gigabit Ethernet, 4-port Gigabit Ethernet, and 4-port Fast Ethernet PICs only.
Options	mac-address —Configure the MAC address of the remote site. inet-address —Configure the IP address of the remote site.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Ethernet TCC</i> • <i>Example: Configuring an Ethernet TCC or Extended VLAN TCC</i> • proxy on page 159 • <i>Junos OS VPNs Library for Routing Devices</i>

rpf-check (Dynamic Profiles)

Syntax	<pre>rpf-check { fail-filter <i>filter-name</i>; mode loose; }</pre>
Hierarchy Level	[edit dynamic-profiles <i>profile-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced in Junos OS Release 9.6.
Description	<p>Check whether traffic is arriving on an expected path. You can include this statement with the inet protocol family only.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Unicast RPF Strict Mode on page 54• <i>Configuring Unicast RPF and Fail Filters in Dynamic Profiles for Subscriber Interfaces</i>

rpf-check (interfaces)

Syntax	<pre>rpf-check { fail-filter <i>filter-name</i>; mode loose; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Check whether traffic is arriving on an expected path. You can include this statement with the inet or inet6 protocol family only.</p> <p>The mode statement is explained separately.</p>
Options	fail-filter —A filter to evaluate when packets are received on the interface. If the RPF check fails, this optional filter is evaluated. If the fail filter is not configured, the default action is to silently discard the packet.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Unicast RPF Strict Mode on page 54 • Configuring Unicast RPF Loose Mode on page 56 • Example: Configuring Unicast Reverse-Path-Forwarding Check on page 64

rtvbr

Syntax	<code>rtvbr peak rate sustained rate burst length;</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> atm-options vpi <i>vpi-identifier</i> shaping]</code> , <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i> shaping]</code> , <code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> shaping]</code> , <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> shaping]</code> , <code>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i> shaping]</code>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>For ATM2 IQ PICs only, define the real-time variable bandwidth utilization in the traffic-shaping profile.</p> <p>When you configure the real-time bandwidth utilization, you must specify all three options (burst, peak, and sustained). You can specify the rate in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000). You can also specify the rate in cells per second by entering a decimal number followed by the abbreviation c; values expressed in cells per second are converted to bits per second using the formula 1 cps = 384 bps.</p>
Default	If the rtvbr statement is not included, bandwidth utilization is unlimited.
Options	<p>burst length—Burst length, in cells. If you set the length to 1, the peak traffic rate is used. Range: 1 through 4000 cells</p> <p>peak rate—Peak rate, in bits per second or cells per second. Range: For ATM2 IQ OC3 and OC12 interfaces, 33 Kbps through 542,526,792 bps. For ATM2 IQ OC48 interfaces, 33 Kbps through 2,170,107,168 bps. For ATM2 IQ DS3 and E3 interfaces, 33 Kbps through the maximum rate, which depends on the ATM encapsulation and framing you configure..</p> <p>sustained rate—Sustained rate, in bps or cps. Range: For ATM2 IQ OC3 and OC12 interfaces, 33 Kbps through 542,526,792 bps. For ATM2 IQ OC48 interfaces, 33 Kbps through 2,170,107,168 bps. For ATM2 IQ DS3 and E3 interfaces, from 33 Kbps through the maximum rate, which depends on the ATM encapsulation and framing you configure.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Configuring ATM CBR</i>• <i>Configuring ATM2 IQ Real-Time VBR</i>• <i>Applying Scheduler Maps to Logical ATM Interfaces</i>

- [cbr on page 111](#)
- [vbr on page 187](#)

sampling (Interfaces)

Syntax	sampling <i>direction</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure the direction of traffic to be sampled.
Options	<p><i>direction</i> can be one of the following:</p> <p>input—Configure at least one expected ingress point.</p> <p>output—Configure at least one expected egress point.</p> <p>input output—On a single interface, configure at least one expected ingress point and one expect egress point.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Junos OS Services Interfaces Library for Routing Devices</i> • <i>Configuring Flow Monitoring</i>

service (Logical Interfaces)

Syntax	<pre>service { input { service-set service-set-name <service-filter filter-name>; post-service-filter filter-name; } output { service-set service-set-name <service-filter filter-name>; } }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Define one or more service sets and filters, and one postservice filter to be applied to an interface.
Options	The remaining statements are explained separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Junos OS Services Interfaces Library for Routing Devices</i>

service-filter (Interfaces)

Syntax	<code>service-filter <i>filter-name</i>;</code>
Hierarchy Level	[edit <code>interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>inet</i> service (<i>input</i> <i>output</i>) service-set <i>service-set-name</i></code>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>inet</i> service (<i>input</i> <i>output</i>) service-set <i>service-set-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Define the filter to be applied to traffic before it is accepted for service processing. Configuration of a service filter is optional; if you include the service-set statement without a service-filter definition, Junos OS assumes the match condition is true and selects the service set for processing automatically.
Options	<i>filter-name</i> —Identifies the filter to be applied in service processing.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Applying Filters and Services to Interfaces</i> • <i>Junos OS Services Interfaces Library for Routing Devices</i>

service-set

Syntax	<code>service-set <i>service-set-name</i>;</code>
Hierarchy Level	[edit <code>interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>inet</i> service (<i>input</i> <i>output</i>) service-set <i>service-set-name</i></code>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>inet</i> service (<i>input</i> <i>output</i>) service-set <i>service-set-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Define one or more service sets to be applied to an interface. If you define multiple service sets, the Junos OS evaluates the filters in the order in which they appear in the configuration.
Options	<i>service-set-name</i> —Identifies the service set.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Junos OS Services Interfaces Library for Routing Devices</i>

shaping

Syntax	<pre>shaping { (cbr <i>rate</i> rtvbr <i>peak rate</i> <i>sustained rate</i> <i>burst length</i> vbr <i>peak rate</i> <i>sustained rate</i> <i>burst length</i>); queue-length <i>number</i>; }</pre>
Hierarchy Level	<pre>[edit interfaces <i>interface-name</i> atm-options vpi <i>vpi-identifier</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i>]</pre>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>For ATM encapsulation only, define the traffic-shaping profile.</p> <p>For Circuit Emulation PICs, specify traffic shaping in the ingress and egress directions.</p> <p>For ATM2 IQ interfaces, changing or deleting VP tunnel traffic shaping causes all logical interfaces on a VP to be deleted and then re-added.</p> <p>VP tunnels are not supported on multipoint interfaces.</p> <p>The statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Defining Virtual Path Tunnels</i>• <i>Defining the ATM Traffic-Shaping Profile</i>• <i>Configuring ATM QoS or Shaping</i>• <i>Applying Scheduler Maps to Logical ATM Interfaces</i>

source-class-usage

Syntax	source-class-usage { <i>direction</i> ; }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet accounting], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet accounting], [edit routing-instances <i>routing-instance-name</i> vrf-table-label]
Release Information	Statement introduced before Junos OS Release 7.4. Support for the vrf-table-label statement added in Junos OS Release 9.3.
Description	Enable packet counters on an interface that count packets that arrive from specific prefixes on the provider core router and are destined for specific prefixes on the customer edge router.
Options	<i>direction</i> can be one of the following: input —Configure at least one expected ingress point. output —Configure at least one expected egress point. input output —On a single interface, configure at least one expected ingress point and one expect egress point.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Enabling Source Class and Destination Class Usage on page 25 • accounting on page 105 • destination-class-usage on page 115 • <i>Junos OS Services Interfaces Library for Routing Devices</i> • <i>vrf-table-label</i>

targeted-broadcast

Syntax	targeted-broadcast { forward-and-send-to-re; forward-only; }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced in Junos OS Release 10.2.
Description	<p>Specify the IP packets destined for a Layer 3 broadcast address to be forwarded to both an egress interface and the Routing Engine, or to an egress interface only. The packets are broadcast only if the egress interface is a LAN interface.</p> <p>The statements are explained separately.</p>
Default	When this statement is not included, broadcast packets are sent to the Routing Engine only.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Targeted Broadcast on page 61• Understanding Targeted Broadcast on page 63

then

Syntax	then { discard; }
Hierarchy Level	[edit firewall hierarchical-policer aggregate], [edit firewall hierarchical-policer premium]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	On M40e, M120, and M320 (with FFPC and SFPC) edge routers and T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs, discard packets when a specified bandwidth or burst limits for an aggregate level of a hierarchical policer is reached.
Options	discard —Discard packets if condition is met.
Required Privilege Level	firewall—To view this statement in the configuration. firewall-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Applying Policers on page 11 • <i>Junos OS Class of Service Library for Routing Devices</i>

translate-discard-eligible

Syntax	(translate-discard-eligible no-translate-discard-eligible);
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>ccc</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>ccc</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For interfaces with encapsulation type Frame Relay CCC, enable or disable translation of Frame Relay discard eligible (DE) control bits.
Default	DE bit translation is disabled.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Frame Relay Control Bit Translation</i>

translate-fecn-and-becn

Syntax	(translate-fecn-and-becn no-translate-fecn-and-becn);
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>ccc</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>ccc</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For interfaces with encapsulation type Frame Relay CCC, enable or disable translation of Frame Relay forward explicit congestion notification (FECN) control bits and Frame Relay backward explicit congestion notification (BECN) control bits.
Default	FECN and BECN bit translation is disabled.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Frame Relay Control Bit Translation</i>

unit (Dynamic Profiles Standard Interface)

```

Syntax  unit logical-unit-number {
        auto-configure {
            agent-circuit-identifier {
                dynamic-profile profile-name;
            }
        }
        dial-options {
            ipsec-interface-id name;
            l2tp-interface-id name;
            (shared | dedicated);
        }
        encapsulation (atm-ccc-cell-relay | atm-ccc-vc-mux | atm-cisco-nlpid | atm-tcc-vc-mux
            | atm-mlppp-llc | atm-nlpid | atm-ppp-llc | atm-ppp-vc-mux | atm-snap | atm-tcc-snap
            | atm-vc-mux | ether-over-atm-llc | ether-vpls-over-atm-llc | ether-vpls-over-fr |
            ether-vpls-over-ppp | ethernet | frame-relay-ccc | frame-relay-ppp | frame-relay-tcc |
            frame-relay-ether-type | frame-relay-ether-type-tcc | multilink-frame-relay-end-to-end
            | multilink-ppp | ppp-over-ether | ppp-over-ether-over-atm-llc | vlan-bridge | vlan-ccc |
            vlan-vci-ccc | vlan-tcc | vlan-vpls);
        family family {
            access-concentrator name;
            address address;
            direct-connect;
            duplicate-protection;
            dynamic-profile profile-name;
            filter {
                adf {
                    counter;
                    input-precedence precedence;
                    not-mandatory;
                    output-precedence precedence;
                    rule rule-value;
                }
                input filter-name (
                    precedence precedence;
                )
                output filter-name {
                    precedence precedence;
                }
            }
            max-sessions number;
            max-sessions-vs-a-ignore;
            rpf-check {
                fail-filter filter-name;
                mode loose;
            }
            service {
                input {
                    service-set service-set-name {
                        service-filter filter-name;
                    }
                }
                post-service-filter filter-name;
            }
        }
    }

```

```
input-vlan-map {
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    (push | swap);
    tag-protocol-id tpid;
    vlan-id number;
}
output {
    service-set service-set-name {
        service-filter filter-name;
    }
}
output-vlan-map {
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    (pop | swap);
    tag-protocol-id tpid;
    vlan-id number;
}
}
service-name-table table-name
short-cycle-protection <lockout-time-min minimum-seconds lockout-time-max
    maximum-seconds>;
unnumbered-address interface-name <preferred-source-address address>;
filter {
    input filter-name;
    output filter-name;
}
keepalives {
    interval seconds;
}
ppp-options {
    chap;
    pap;
}
vlan-id number;
vlan-tags outer [tpid].vlan-id [inner [tpid].vlan-id];
}
}
```

Hierarchy Level [edit [dynamic-profiles](#) *profile-name* [interfaces](#) *interface-name*]

Release Information Statement introduced in Junos OS Release 9.2.

Description Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options *logical-unit-number*—The specific unit number of the interface you want to assign to the dynamic profile, or one of the following Junos OS predefined variables:

- **\$junos-underlying-interface-unit**—For static VLANs, the unit number variable. The static unit number variable is dynamically replaced with the client unit number when the client session begins. The client unit number is specified by the DHCP when it accesses the subscriber network.
- **\$junos-interface-unit**—The unit number variable on a dynamic underlying VLAN interface for which you want to enable the creation of dynamic VLAN subscriber interfaces based on agent circuit identifier information.

The remaining statements are explained separately.

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

Related Documentation

- *Configuring Dynamic Underlying VLAN Interfaces to Use Agent Circuit Identifier Information*
- *Configuring Static Underlying VLAN Interfaces to Use Agent Circuit Identifier Information*
- *Agent Circuit Identifier-Based Dynamic VLANs Components Overview*

unit

```
Syntax  unit logical-unit-number {
        accept-source-mac {
            mac-address mac-address {
                policer {
                    input cos-policer-name;
                    output cos-policer-name;
                }
            }
        }
        accounting-profile name;
        advisory-options {
            downstream-rate rate;
            upstream-rate rate;
        }
        allow-any-vci;
        atm-scheduler-map (map-name | default);
        backup-options {
            interface interface-name;
        }
        bandwidth rate;
        cell-bundle-size cells;
        clear-dont-fragment-bit;
        compression {
            rtp {
                maximum-contexts number <force>;
                f-max-period number;
                queues [queue-numbers];
                port {
                    minimum port-number;
                    maximum port-number;
                }
            }
        }
        compression-device interface-name;
        copy-tos-to-outer-ip-header;
        demux-destination family;
        demux-source family;
        demux-options {
            underlying-interface interface-name;
        }
        description text;
        interface {
            l2tp-interface-id name;
            (dedicated | shared);
        }
        dialer-options {
            activation-delay seconds;
            callback;
            callback-wait-period time;
            deactivation-delay seconds;
            dial-string [dial-string-numbers];
            idle-timeout seconds;
```

```

incoming-map {
  caller caller-id | accept-all;
  initial-route-check seconds;
  load-interval seconds;
  load-threshold percent;
  pool pool-name;
  redial-delay time;
  watch-list {
    [routes];
  }
}
}
disable;
disable-mlppp-inner-ppp-pfc;
dlci dlci-identifier;
drop-timeout milliseconds;
dynamic-call-admission-control {
  activation-priority priority;
  bearer-bandwidth-limit kilobits-per-second;
}
encapsulation type;
epd-threshold cells plp1 cells;
family family-name {
  ... the family subhierarchy appears after the main [edit interfaces interface-name unit
    logical-unit-number] hierarchy ...
}
fragment-threshold bytes;
inner-vlan-id-range start start-id end end-id;
input-vlan-map {
  (pop | pop-pop | pop-swap | push | push-push | swap |
  swap-push | swap-swap);
  inner-tag-protocol-id tpid;
  inner-vlan-id number;
  tag-protocol-id tpid;
  vlan-id number;
}
interleave-fragments;
inverse-arp;
layer2-policer {
  input-policer policer-name;
  input-three-color policer-name;
  output-policer policer-name;
  output-three-color policer-name;
}
link-layer-overhead percent;
minimum-links number;
mrru bytes;
multicast-dlci dlci-identifier;
multicast-vci vpi-identifier.vci-identifier;
multilink-max-classes number;
multipoint;
oam-liveness {
  up-count cells;
  down-count cells;
}
oam-period (disable | seconds);

```

```
output-vlan-map {
  (pop | pop-pop | pop-swap | push | push-push | swap |
  swap-push | swap-swap);
  inner-tag-protocol-id tpid;
  inner-vlan-id number;
  tag-protocol-id tpid;
  vlan-id number;
}
passive-monitor-mode;
peer-unit unit-number;
plp-to-clp;
point-to-point;
ppp-options {
  chap {
    access-profile name;
    default-chap-secret name;
    local-name name;
    passive;
  }
  compression {
    acfc;
    pfc;
  }
  dynamic-profile profile-name;
  lcp-restart-timer milliseconds;
  loopback-clear-timer seconds;
  ncp-restart-timer milliseconds;
  pap {
    access-profile name;
    default-pap-password password;
    local-name name;
    local-password password;
    passive;
  }
}
pppoe-options {
  access-concentrator name;
  auto-reconnect seconds;
  (client | server);
  service-name name;
  underlying-interface interface-name;
}
pppoe-underlying-options {
  access-concentrator name;
  direct-connect;
  dynamic-profile profile-name;
  max-sessions number;
}
proxy-arp;
service-domain (inside | outside);
shaping {
  (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate burst length);
  queue-length number;
}
short-sequence;
```

```

targeted-distribution;
transmit-weight number;
(traps | no-traps);
trunk-bandwidth rate;
trunk-id number;
tunnel {
    backup-destination address;
    destination address;
    key number;
    routing-instance {
        destination routing-instance-name;
    }
    source source-address;
    ttl number;
}
vci vpi-identifier.vci-identifier;
vci-range start start-vci end end-vci;
vpi vpi-identifier;
vlan-id number;
vlan-id-range number-number;
vlan-tags inner tpid.vlan-id outer tpid.vlan-id;
family family {
    accounting {
        destination-class-usage;
        source-class-usage {
            (input | output | input output);
        }
    }
    access-concentrator name;
    address address {
        ... the address subhierarchy appears after the main [edit interfaces interface-name unit
        logical-unit-number family family-name] hierarchy ...
    }
    bundle interface-name;
    core-facing;
    demux-destination {
        destination-prefix;
    }
    demux-source {
        source-prefix;
    }
    direct-connect;
    duplicate-protection;
    dynamic-profile profile-name;
    filter {
        group filter-group-number;
        input filter-name;
        input-list [filter-names];
        output filter-name;
        output-list [filter-names];
    }
    interface-mode (access | trunk);
    ipsec-sa sa-name;
    keep-address-and-control;
    mac-validate (loose | strict);
    max-sessions number;

```

```
mtu bytes;
multicast-only;
no-redirects;
policer {
    arp policer-template-name;
    input policer-template-name;
    output policer-template-name;
}
primary;
protocols [inet iso mpls];
proxy inet-address address;
receive-options-packets;
receive-ttl-exceeded;
remote (inet-address address | mac-address address);
rpf-check {
    fail-filter filter-name
    mode loose;
}
sampling {
    input;
    output;
}
service {
    input {
        post-service-filter filter-name;
        service-set service-set-name <service-filter filter-name>;
    }
    output {
        service-set service-set-name <service-filter filter-name>;
    }
}
service-name-table table-name
(translate-discard-eligible | no-translate-discard-eligible);
(translate-fecn-and-becn | no-translate-fecn-and-becn);
translate-plp-control-word-de;
unnumbered-address interface-name destination address
    destination-profile profile-name;
vlan-id number;
vlan-id-list [number number-number];
address address {
    arp ip-address (mac | multicast-mac) mac-address <publish>;
    broadcast address;
    destination address;
    destination-profile name;
    eui-64;
    master-only;
    multipoint-destination address {
        dlci dlci-identifier;
        epd-threshold cells <plp cells>;
        inverse-arp;
        oam-liveness {
            up-count cells;
            down-count cells;
        }
        oam-period (disable | seconds);
        shaping {
```

```

        (cbr rate | rtvbr burst length peak rate sustained rate | vbr burst length peak rate
         sustained rate);
        queue-length number;
    }
    vci vpi-identifier.vci-identifier;
}
preferred;
primary;
(vrrp-group | vrrp-inet6-group) group-number {
    (accept-data | no-accept-data);
    advertise-interval seconds;
    authentication-type authentication;
    authentication-key key;
    fast-interval milliseconds;
    (preempt | no-preempt) {
        hold-time seconds;
    }
    priority number;
    track {
        interface interface-name {
            bandwidth-threshold bits-per-second priority-cost number;
        }
        priority-hold-time seconds;
        route ip-address/prefix-length routing-instance instance-name priority-cost cost;
    }
    virtual-address [addresses];
    virtual-link-local-address ipv6-address;
    vrrp-inherit-from {
        active-interface interface-name;
        active-group group-number;
    }
}
}
}
}

```

Hierarchy Level [edit interfaces *interface-name*],
 [edit logical-systems *logical-system-name* interfaces *interface-name*],
 [edit interfaces interface-set *interface-set-name* interface *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options *logical-unit-number*—Number of the logical unit.
Range: 0 through 1,073,741,823 for demux and PPPoE static interfaces. 0 through 16,385 for all other static interface types.

The remaining statements are explained separately.

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

- Related Documentation**
- [Configuring Logical Interface Properties](#)
 - [Junos OS Services Interfaces Library for Routing Devices](#)

unnumbered-address (Demux)

- Syntax** `unnumbered-address interface-name <preferred-source-address address>;`
- Hierarchy Level** [edit [interfaces](#) *interface-name* [unit](#) *logical-unit-number* [family](#) inet],
[edit [logical-systems](#) *logical-system-name* [interfaces](#) *interface-name* [unit](#) *logical-unit-number* [family](#) inet]
- Release Information** Statement introduced in Junos OS Release 8.2.
preferred-source-address option introduced in Junos OS Release 9.0.
IP demultiplexing interfaces supported in Junos OS Release 9.2.
- Description** For IP demultiplexing interfaces, enable the local address to be derived from the specified interface. Configuring an unnumbered interface enables IP processing on the interface without assigning an explicit IP address to the interface.
- Options** *interface-name*—Name of the interface from which the local address is derived. The specified interface must have a logical unit number and a configured IP address, and must not be an unnumbered interface.
- The **preferred-source-address** statement is explained separately.
- Required Privilege Level** interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.
- Related Documentation**
- [Configuring an Unnumbered Interface on page 44](#)
 - [address on page 102](#)
 - [Junos System Basics Configuration Guide](#)

unnumbered-address (Dynamic Profiles)

Syntax	<code>unnumbered-address interface-name <preferred-source-address address>;</code>
Hierarchy Level	<p>[edit dynamic-profiles <i>profile-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>],</p> <p>[edit dynamic-profiles <i>profile-name</i> interfaces demux0 unit <i>logical-unit-number</i> family <i>family</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 9.2.</p> <p>\$junos-preferred-source-address variable support added in Junos OS Release 9.6.</p> <p>Support for the \$junos-loopback-interface predefined variable introduced in Junos OS Release 9.6.</p>
Description	<p>For Ethernet interfaces, enable the local address to be derived from the specified interface. Configuring unnumbered Ethernet interfaces enables IP processing on the interface without assigning an explicit IP address to the interface. To configure unnumbered address dynamically, include the \$junos-loopback-interface-address predefined variable.</p> <p>You can configure unnumbered address support on Ethernet interfaces for IPv4 and IPv6 address families.</p>
Options	<p>interface-name—Name of the interface from which the local address is derived. The specified interface must have a logical unit number, a configured IP address, and must not be an unnumbered interface. This value can be a specific interface name or the \$junos-loopback-interface dynamic variable.</p> <p>When defining the unnumbered-address statement using a static interface, keep the following in mind:</p> <ul style="list-style-type: none"> If you choose to include the routing-instance statement at the [edit dynamic-profiles] hierarchy level, that statement must be configured with a valid, static routing instance value. In addition, whatever static unnumbered interface you specify must belong to that routing instance. If you choose to not include the routing-instance statement at the [edit dynamic-profiles] hierarchy level, the unnumbered-address statement uses the default routing instance. The use of the default routing instance requires that the unnumbered interface be configured statically and that it reside in the default routing instance. <p>When defining the unnumbered-address statement using the \$junos-loopback-interface dynamic variable, keep the following in mind:</p> <ul style="list-style-type: none"> To use the \$junos-loopback-interface dynamic variable, the dynamic profile must also contain the routing-instance statement configured with the \$junos-routing-instance dynamic variable at the [edit dynamic-profiles] hierarchy level. The applied loopback interface is based on the dynamically obtained routing instance of the subscriber.

address—(Optional) Secondary IP address of the donor interface. Configuring the preferred source address enables you to use an IP address other than the primary IP address on some of the unnumbered Ethernet interfaces in your network. This value can be a static IP address, the **\$junos-preferred-source-address** dynamic variable for the inet family, or **\$junos-preferred-source-ipv6-address** dynamic variable for the inet6 family.

When defining the **preferred-source-address** value using a static IP address, keep the following in mind:

- The unnumbered interface must be statically configured.
- The IP address specified as the **preferred-source-address** must be configured in the specified unnumbered interface.

When defining the **preferred-source-address** value using the **\$junos-preferred-source-address** or **\$junos-preferred-source-ipv6-address** dynamic variables, keep the following in mind:

- You must configure the **unnumbered-address** statement using the **\$junos-loopback-interface** dynamic variable.
- You must configure the **routing-instance** statement using the **\$junos-routing-instance** dynamic variable at the **[edit dynamic-profiles]** hierarchy level.
- The preferred source address chosen is based on the dynamically applied loopback address which is in turn derived from the dynamically obtained routing instance of the subscriber. The configured loopback address with the closest network match to the user IP address is selected as the preferred source address.

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

Related Documentation	<ul style="list-style-type: none">• Configuring an Unnumbered Interface on page 44• <i>Dynamic Profiles Overview</i>
------------------------------	---

unnumbered-address (Ethernet)

Syntax	<code>unnumbered-address interface-name <preferred-source-address address>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced in Junos OS Release 8.2. preferred-source-address option introduced in Junos OS Release 9.0.
Description	For Ethernet interfaces, enable the local address to be derived from the specified interface. Configuring an unnumbered Ethernet interface enables IP processing on the interface without assigning an explicit IP address to the interface.
Options	interface-name —Name of the interface from which the local address is derived. The specified interface must have a logical unit number and a configured IP address, and must not be an unnumbered interface. The preferred-source-address statement is explained separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring an Unnumbered Interface on page 44 • address on page 102 • <i>Junos System Basics Configuration Guide</i>

unnumbered-address (PPP)

Syntax	<code>unnumbered-address interface-name destination address destination-profile profile-name;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	For interfaces with PPP encapsulation, enable the local address to be derived from the specified interface.
Options	<i>interface-name</i> —Interface from which the local address is derived. The interface name must include a logical unit number and must have a configured address. The remaining statements are explained separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring IPCP Options on page 42• address on page 102• negotiate-address on page 150• <i>Junos OS Administration Library for Routing Devices</i>

vbr

Syntax	<code>vbr peak <i>rate</i> sustained <i>rate</i> burst <i>length</i>;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> atm-options vpi <i>vpi-identifier</i> shaping],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i> shaping],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> shaping],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i> shaping],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> shaping]</p>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>For ATM encapsulation only, define the variable bandwidth utilization in the traffic-shaping profile.</p> <p>When you configure the variable bandwidth utilization, you must specify all three options (burst, peak, and sustained). You can specify the rate in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000). You can also specify the rate in cells per second by entering a decimal number followed by the abbreviation c; values expressed in cells per second are converted to bits per second by means of the formula 1 cps = 384 bps.</p>
Default	If the vbr statement is not specified, bandwidth utilization is unlimited.
Options	<p>burst <i>length</i>—Burst length, in cells. If you set the length to 1, the peak traffic rate is used. Range: 1 through 4000 cells</p> <p>peak <i>rate</i>—Peak rate, in bits per second or cells per second. Range: For ATM1 interfaces, 33 Kbps through 135.6 Mbps (ATM OC3); 33 Kbps through 276 Mbps (ATM OC12). For ATM2 IQ OC3 and OC12 interfaces, 33 Kbps through 542,526,792 bps. For ATM2 IQ OC48 interfaces, 33 Kbps through 2,170,107,168 bps. For ATM2 IQ DS3 and E3 interfaces, from 33 Kbps through the maximum rate, which depends on the ATM encapsulation and framing you configure.</p> <p>sustained <i>rate</i>—Sustained rate, in bits per second or cells per second. Range: For ATM1 interfaces, 33 Kbps through 135.6 Mbps (ATM OC3); 33 Kbps through 276 Mbps (ATM OC12). For ATM2 IQ OC3 and OC12 interfaces, 33 Kbps through 542,526,792 bps. For ATM2 IQ OC48 interfaces, 33 Kbps through 2,170,107,168 bps. For ATM2 IQ DS3 and E3 interfaces, from 33 Kbps through the maximum rate, which depends on the ATM encapsulation and framing you configure.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring ATM CBR</i> • <i>Applying Scheduler Maps to Logical ATM Interfaces</i>

- [cbr on page 111](#)
- [rtvbr on page 164](#)
- [shaping on page 168](#)

vci

Syntax	<code>vci vpi-identifier.vci-identifier;</code>
Hierarchy Level	<code>[edit interfaces at-fpc/pic/port unit logical-unit-number],</code> <code>[edit interfaces at-fpc/pic/port unit logical-unit-number family family address address</code> <code> multipoint-destination address],</code> <code>[edit logical-systems logical-system-name interfaces at-fpc/pic/port unit logical-unit-number],</code> <code>[edit logical-systems logical-system-name interfaces at-fpc/pic/port unit logical-unit-number</code> <code> family family address address multipoint-destination address]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 12.2 for the ACX Series Universal Access routers.
Description	<p>For ATM point-to-point logical interfaces only, configure the virtual circuit identifier (VCI) and virtual path identifier (VPI).</p> <p>To configure a VPI for a point-to-multipoint interface, specify the VPI in the multipoint-destination statement.</p> <p>VCIs 0 through 31 are reserved for specific ATM values designated by the ATM Forum.</p>
Options	<p>vci-identifier—ATM virtual circuit identifier. Unless you configure the interface to use promiscuous mode, this value cannot exceed the highest-numbered VC configured for the interface with the maximum-vcs option of the vpi statement.</p> <p>Range: 0 through 4089 or 0 through 65,535 with promiscuous mode, with VCIs 0 through 31 reserved.</p> <p>vpi-identifier—ATM virtual path identifier.</p> <p>Range: 0 through 255</p> <p>Default: 0</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Point-to-Point ATM1 or ATM2 IQ Connection• Applying Scheduler Maps to Logical ATM Interfaces• multipoint-destination on page 149• promiscuous-mode• vpi (ATM CCC Cell-Relay Promiscuous Mode)

vlan-id (Logical Port in Bridge Domain)

Syntax	<code>vlan-id <i>number</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge]
Release Information	Statement introduced in Junos OS Release 9.2.
Description	The VLAN ID configured on the logical port. Received packets with no VLAN tags are forwarded within the bridge domain with the matching VLAN ID.
Options	number —The VLAN ID. Range: 1 through 4095
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring a Logical Interface for Access Mode</i>

vlan-id-list (Interface in Bridge Domain)

Syntax	<code>vlan-id-list [<i>number number-number</i>];</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge]
Release Information	Statement introduced in Junos OS Release 9.2.
Description	Configure a logical interface to forward packets and learn MAC addresses within each bridge domain configured with a VLAN ID that matches a VLAN ID specified in the list. VLAN IDs can be entered individually using a space to separate each ID, entered as an inclusive list separating the starting VLAN ID and ending VLAN ID with a hyphen, or a combination of both.
Options	<i>number number</i> —Individual VLAN IDs separated by a space. <i>number-number</i> —Starting VLAN ID and ending VLAN ID in an inclusive range. Range: 1 through 4095
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring a Logical Interface for Trunk Mode</i> • <i>Configuring the VLAN ID List for a Trunk Interface</i>

PART 3

Administration

- [Monitoring Commands on page 193](#)
- [ANCP Operational Mode Commands on page 259](#)
- [BFD Operational Mode Commands on page 295](#)
- [BGP Operational Mode Commands on page 307](#)
- [ES-IS Operational Mode Commands on page 349](#)
- [IP Multicast Operational Mode Commands on page 359](#)
- [IPv6 Operational Mode Commands on page 589](#)
- [IS-IS Operational Mode Commands on page 597](#)
- [LLDP Operational Mode Commands on page 655](#)
- [MVRP Operational Mode Commands on page 673](#)
- [OSPF Operational Mode Commands on page 685](#)
- [Protocol-Independent Routing Operational Mode Commands on page 763](#)
- [RIP Operational Mode Commands on page 969](#)
- [RIPng Operational Mode Commands on page 979](#)
- [Firewall Filter Operational Mode Commands on page 987](#)
- [Layer 2 Bridging and Switching Operational Mode Commands on page 1009](#)
- [VPN Operational Mode Commands on page 1037](#)

CHAPTER 5

Monitoring Commands

- `show interfaces (10-Gigabit Ethernet)`
- `show interfaces (Fast Ethernet)`
- `show interfaces (Gigabit Ethernet)`

show interfaces (10-Gigabit Ethernet)

Syntax	<code>show interfaces <i>xe-fpc/pic/port</i></code> <code><brief detail extensive terse></code> <code><descriptions></code> <code><media></code> <code><snmp-index <i>snmp-index</i>></code> <code><statistics></code>
Release Information	Command introduced in Junos OS Release 8.0.
Description	(M320, M120, MX Series, and T Series routers and EX Series switches only) Display status information about the specified 10-Gigabit Ethernet interface.
Options	<p><code><i>xe-fpc/pic/port</i></code>—Display standard information about the specified 10-Gigabit Ethernet interface.</p> <p><code>brief detail extensive terse</code>—(Optional) Display the specified level of output.</p> <p><code>descriptions</code>—(Optional) Display interface description strings.</p> <p><code>media</code>—(Optional) Display media-specific information about network interfaces.</p> <p><code>snmp-index <i>snmp-index</i></code>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><code>statistics</code>—(Optional) Display static interface statistics.</p>
Required Privilege Level	view
List of Sample Output	<p>show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, IQ2) on page 209</p> <p>show interfaces extensive (10-Gigabit Ethernet, WAN PHY Mode) on page 212</p> <p>show interfaces extensive (10-Gigabit Ethernet, DWDM OTN PIC) on page 214</p> <p>show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode) on page 216</p> <p>show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Transmit-Only) on page 216</p> <p>show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only) on page 217</p>
Output Fields	See Table 3 on page 195 for the output fields for the show interfaces (10-Gigabit Ethernet) command.

Table 3: show interfaces Gigabit Ethernet Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels
Loopback	Loopback status: Enabled or Disabled . If loopback is enabled, type of loopback: Local or Remote .	All levels
Source filtering	Source filtering status: Enabled or Disabled .	All levels
LAN-PHY mode	10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.	All levels
WAN-PHY mode	10-Gigabit Ethernet interface operating in Wide Area Network Physical Layer Device (WAN PHY) mode. WAN PHY allows 10-Gigabit Ethernet wide area links to use fiber-optic cables and other devices intended for SONET/SDH.	All levels
Unidirectional	Unidirectional link mode status for 10-Gigabit Ethernet interface: Enabled or Disabled for parent interface; Rx-only or Tx-only for child interfaces.	All levels
Flow control	Flow control status: Enabled or Disabled .	All levels
Auto-negotiation	(Gigabit Ethernet interfaces) Autonegotiation status: Enabled or Disabled .	All levels
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: <ul style="list-style-type: none"> • Online—Autonegotiation is manually configured as online. • Offline—Autonegotiation is manually configured as offline. 	All levels
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> .	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .	All levels

Table 3: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output	
Link flags	Information about the link. Possible values are described in the “Links Flags” section under <i>Common Output Fields Description</i> .	All levels	
Wavelength	(10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).	All levels	
Frequency	(10-Gigabit Ethernet DWDM interfaces only) Displays the frequency associated with the configured wavelength, in terahertz (THz).	All levels	
CoS queues	Number of CoS queues configured.	detail extensive none	
Schedulers	(Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces only) Number of CoS schedulers configured.	extensive	
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	detail extensive	
Current address	Configured MAC address.	detail extensive none	
Hardware address	Hardware MAC address.	detail extensive none	
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .	detail extensive none	
Input Rate	Input rate in bits per second (bps) and packets per second (pps). The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	None specified	
Output Rate	Output rate in bps and pps. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	None specified	
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive	
Egress account overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for egress traffic.	detail extensive	
Ingress account overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for ingress traffic.	detail extensive	detail extensive

Table 3: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Output bytes—Number of bytes transmitted on the interface. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. <p>Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes.</p> <p>For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 3 on page 195.</p>	detail extensive
Input errors	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the ignore-l3-incompletes statement. • L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame. • L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • FIFO errors—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • Resource errors—Sum of transmit drops. 	extensive

Table 3: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Output errors	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Collisions—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • FIFO errors—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	extensive
Egress queues	Total number of egress queues supported on the specified interface.	detail extensive
Queue counters (Egress)	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	detail extensive
Ingress queues	Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.	extensive
Queue counters (Ingress)	<p>CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	extensive

Table 3: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Active alarms and Active defects	<p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the routing device configuration, an alarm can ring the red or yellow alarm bell on the routing device, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value None or Link.</p> <ul style="list-style-type: none"> • None—There are no active defects or alarms. • Link—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning. 	detail extensive none
OTN alarms	Active OTN alarms identified on the interface.	detail extensive
OTN defects	OTN defects received on the interface.	detail extensive
OTN FEC Mode	<p>The FECmode configured on the interface.</p> <ul style="list-style-type: none"> • efec—Enhanced forward error correction (EFEC) is configured to detect and correct bit errors. • gfec—G.709 Forward error correction (GFEC) mode is configured to detect and correct bit errors. • none—FEC mode is not configured. 	detail extensive
OTN Rate	<p>OTN mode.</p> <ul style="list-style-type: none"> • fixed-stuff-bytes—Fixed stuff bytes 11.0957 Gbps. • no-fixed-stuff-bytes—No fixed stuff bytes 11.0491 Gbps. • pass-through—Enable OTN passthrough mode. • no-pass-through—Do not enable OTN passthrough mode. 	detail extensive
OTN Line Loopback	Status of the line loopback, if configured for the DWDM OTN PIC. Its value can be: enabled or disabled .	detail extensive
OTN FEC statistics	<p>The forward error correction (FEC) counters for the DWDM OTN PIC.</p> <ul style="list-style-type: none"> • Corrected Errors—The count of corrected errors in the last second. • Corrected Error Ratio—The corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits. 	detail extensive
OTN FEC alarms	<p>OTN FEC excessive or degraded error alarms triggered on the interface.</p> <ul style="list-style-type: none"> • FEC Degrade—OTU FEC Degrade defect. • FEC Excessive—OTU FEC Excessive Error defect. 	detail extensive
OTN OC	<p>OTN OC defects triggered on the interface.</p> <ul style="list-style-type: none"> • LOS—OC Loss of Signal defect. • LOF—OC Loss of Frame defect. • LOM—OC Loss of Multiframe defect. • Wavelength Lock—OC Wavelength Lock defect. 	detail extensive

Table 3: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
OTN OTU	OTN OTU defects detected on the interface <ul style="list-style-type: none"> • AIS—OTN AIS alarm. • BDI—OTN OTU BDI alarm. • IAE—OTN OTU IAE alarm. • TTIM—OTN OTU TTIM alarm. • SF—OTN ODU bit error rate fault alarm. • SD—OTN ODU bit error rate defect alarm. • TCA-ES—OTN ODU ES threshold alarm. • TCA-SES—OTN ODU SES threshold alarm. • TCA-UAS—OTN ODU UAS threshold alarm. • TCA-BBE—OTN ODU BBE threshold alarm. • BIP—OTN ODU BIP threshold alarm. • BBE—OTN OTU BBE threshold alarm. • ES—OTN OTU ES threshold alarm. • SES—OTN OTU SES threshold alarm. • UAS—OTN OTU UAS threshold alarm. 	detail extensive
Received DAPI	Destination Access Port Interface (DAPI) from which the packets were received.	detail extensive
Received SAPI	Source Access Port Interface (SAPI) from which the packets were received.	detail extensive
Transmitted DAPI	Destination Access Port Interface (DAPI) to which the packets were transmitted.	detail extensive
Transmitted SAPI	Source Access Port Interface (SAPI) to which the packets were transmitted.	detail extensive
PCS statistics	(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device. <ul style="list-style-type: none"> • Bit errors—High bit error rate. Indicates the number of bit errors when the PCS receiver is operating in normal mode. • Errored blocks—Loss of block lock. The number of errored blocks when PCS receiver is operating in normal mode. 	detail extensive

Table 3: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
MAC statistics	<p>Receive and Transmit statistics reported by the PIC's MAC subsystem, including the following:</p> <ul style="list-style-type: none"> • Total octets and total packets—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. For more information, see Table 4 on page 209 • Unicast packets, Broadcast packets, and Multicast packets—Number of unicast, broadcast, and multicast packets. • CRC/Align errors—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • FIFO error—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning. • MAC control frames—Number of MAC control frames. • MAC pause frames—Number of MAC control frames with pause operational code. • Oversized frames—Number of frames that exceed 1518 octets. • Jabber frames—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • Fragment frames—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • VLAN tagged frames—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. • Code violations—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error." 	extensive
OTN Received Overhead Bytes	APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58 Payload Type: 0x08	extensive
OTN Transmitted Overhead Bytes	APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00 Payload Type: 0x08	extensive

Table 3: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Filter statistics	<p>Receive and Transmit statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected.</p> <ul style="list-style-type: none"> • Input packet count—Number of packets received from the MAC hardware that the filter processed. • Input packet rejects—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • Input DA rejects—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the routing device from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local routing device (which the routing device is rejecting). • Input SA rejects—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect. • Output packet count—Number of packets that the filter has given to the MAC hardware. • Output packet pad count—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured. • Output packet error count—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment. • CAM destination filters, CAM source filters—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields should be 0. 	extensive
PMA PHY	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • PHY Lock—Phase-locked loop • PHY Light—Loss of optical signal 	extensive

Table 3: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
WIS section	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • BIP-B1—Bit interleaved parity for SONET section overhead • SEF—Severely errored framing • LOL—Loss of light • LOF—Loss of frame • ES-S—Errored seconds (section) • SES-S—Severely errored seconds (section) • SEFS-S—Severely errored framing seconds (section) 	extensive
WIS line	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. State other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • BIP-B2—Bit interleaved parity for SONET line overhead • REI-L—Remote error indication (near-end line) • RDI-L—Remote defect indication (near-end line) • AIS-L—Alarm indication signal (near-end line) • BERR-SF—Bit error rate fault (signal failure) • BERR-SD—Bit error rate defect (signal degradation) • ES-L—Errored seconds (near-end line) • SES-L—Severely errored seconds (near-end line) • UAS-L—Unavailable seconds (near-end line) • ES-LFE—Errored seconds (far-end line) • SES-LFE—Severely errored seconds (far-end line) • UAS-LFE—Unavailable seconds (far-end line) 	extensive

Table 3: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
WIS path	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • BIP-B3—Bit interleaved parity for SONET section overhead • REI-P—Remote error indication • LOP-P—Loss of pointer (path) • AIS-P—Path alarm indication signal • RDI-P—Path remote defect indication • UNEQ-P—Path unequipped • PLM-P—Path payload label mismatch • ES-P—Errored seconds (near-end STS path) • SES-P—Severely errored seconds (near-end STS path) • UAS-P—Unavailable seconds (near-end STS path) • SES-PFE—Severely errored seconds (far-end STS path) • UAS-PFE—Unavailable seconds (far-end STS path) 	extensive

Table 3: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Autonegotiation information	<p>Information about link autonegotiation.</p> <ul style="list-style-type: none"> • Negotiation status: <ul style="list-style-type: none"> • Incomplete—Ethernet interface has the speed or link mode configured. • No autonegotiation—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation. • Complete—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner status—OK when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner: <ul style="list-style-type: none"> • Link mode—Depending on the capability of the attached Ethernet device, either Full-duplex or Half-duplex. • Flow control—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is None. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), and Symmetric/Asymmetric (link partner supports both PAUSE on receive and transmit or only PAUSE receive). • Remote fault—Remote fault information from the link partner—Failure indicates a receive link error. OK indicates that the link partner is receiving. Negotiation error indicates a negotiation error. Offline indicates that the link partner is going offline. • Local resolution—Information from the link partner: <ul style="list-style-type: none"> • Flow control—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), and Symmetric/Asymmetric (link partner supports both PAUSE on receive and transmit or only PAUSE receive). • Remote fault—Remote fault information. Link OK (no error detected on receive), Offline (local interface is offline), and Link Failure (link error detected on receive). 	extensive
Received path trace, Transmitted path trace	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other router manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the routing device at the other end of the fiber. The transmitted path trace value is the message that this routing device transmits.</p>	extensive
Packet Forwarding Engine configuration	<p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> • Destination slot—FPC slot number. 	extensive

Table 3: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
CoS information	Information about the CoS queue for the physical interface. <ul style="list-style-type: none"> • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Bandwidth bps—Bandwidth allocated to the queue (in bps). • Buffer %—Percentage of buffer space allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available. 	extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under <i>Common Output Fields Description</i> .	All levels

Table 3: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
VLAN-Tag	<p>Rewrite profile applied to incoming or outgoing frames on the outer (Out) VLAN tag or for both the outer and inner (In) VLAN tags.</p> <ul style="list-style-type: none"> push—An outer VLAN tag is pushed in front of the existing VLAN tag. pop—The outer VLAN tag of the incoming frame is removed. swap—The outer VLAN tag of the incoming frame is overwritten with the user specified VLAN tag information. push—An outer VLAN tag is pushed in front of the existing VLAN tag. push-push—Two VLAN tags are pushed in from the incoming frame. swap-push—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame. swap-swap—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user specified VLAN tag value. pop-swap—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame. pop-pop—Both the outer and inner VLAN tags of the incoming frame are removed. 	brief detail extensive none
Demux:	<p>IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following:</p> <ul style="list-style-type: none"> Source Family Inet Destination Family Inet 	detail extensive none
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family. Possible values are described in the “Protocol Field” section under <i>Common Output Fields Description</i> .	detail extensive none
MTU	Maximum transmission unit size on the logical interface.	detail extensive none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	detail extensive none
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> Input bytes, Output bytes—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. Input packets, Output packets—Number of packets received and transmitted on the interface set. 	detail extensive
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	extensive
Local statistics	Number and rate of bytes and packets destined to the routing device.	extensive

Table 3: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Transit statistics	Number and rate of bytes and packets transiting the switch. NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the Output bytes and Output packets interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.	extensive
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive none
Flags	Information about protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	detail extensive
Donor interface	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	detail extensive none
Preferred source address	(Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.	detail extensive none
Input Filters	Names of any input filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.	detail extensive
Output Filters	Names of any output filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.	detail extensive
Mac-Validate Failures	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	detail extensive none
Addresses, Flags	Information about the address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .	detail extensive none
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is inet , the IP address of the interface is also displayed.	brief
Flags	Information about address flag (possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i>).	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interlace.	detail extensive none

Table 3: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

For Gigabit Ethernet IQ PICs, traffic and MAC statistics output varies. [Table 4 on page 209](#) describes the traffic and MAC statistics for two sample interfaces, each of which is sending traffic in packets of 500 bytes (including 478 bytes for the Layer 3 packet, 18 bytes for the Layer 2 VLAN traffic header, and 4 bytes for cyclic redundancy check [CRC] information). In [Table 4 on page 209](#), the **ge-0/3/0** interface is the inbound physical interface, and the **ge-0/0/0** interface is the outbound physical interface. On both interfaces, traffic is carried on logical unit .50 (VLAN 50).

Table 4: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type

Interface Type	Sample Command	Byte and Octet Counts Include	Comments
Inbound physical interface	show interfaces ge-0/3/0 extensive	Traffic statistics: Input bytes: 496 bytes per packet, representing the Layer 2 packet MAC statistics: Received octets: 500 bytes per packet, representing the Layer 2 packet + 4 bytes	The additional 4 bytes are for the CRC.
Inbound logical interface	show interfaces ge-0/3/0.50 extensive	Traffic statistics: Input bytes: 478 bytes per packet, representing the Layer 3 packet	
Outbound physical interface	show interfaces ge-0/0/0 extensive	Traffic statistics: Input bytes: 490 bytes per packet, representing the Layer 3 packet + 12 bytes MAC statistics: Received octets: 478 bytes per packet, representing the Layer 3 packet	For input bytes, the additional 12 bytes includes 6 bytes for the destination MAC address + 4 bytes for VLAN + 2 bytes for the Ethernet type.
Outbound logical interface	show interfaces ge-0/0/0.50 extensive	Traffic statistics: Input bytes: 478 bytes per packet, representing the Layer 3 packet	

Sample Output

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, IQ2)

```

user@host> show interfaces xe-5/0/0 extensive
Physical interface: xe-5/0/0, Enabled, Physical link is Up
  Interface index: 177, SNMP ifIndex: 99, Generation: 178
  Link-level type: Ethernet, MTU: 1518, LAN-PHY mode, Speed: 10Gbps, Loopback:

```

```

None, Source filtering: Enabled,
Flow control: Enabled
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
Link flags     : None
CoS queues    : 8 supported, 4 maximum usable queues
Schedulers    : 1024
Hold-times    : Up 0 ms, Down 0 ms
Current address: 00:14:f6:b9:f1:f6, Hardware address: 00:14:f6:b9:f1:f6
Last flapped  : Never
Statistics last cleared: Never
Traffic statistics:
Input bytes   :          6970332384          0 bps
Output bytes  :              0          0 bps
Input packets :          81050506          0 pps
Output packets:              0          0 pps
IPv6 transit statistics:
Input bytes   :              0
Output bytes  :              0
Input packets :              0
Output packets:              0
Ingress traffic statistics at Packet Forwarding Engine:
Input bytes   :          6970299398          0 bps
Input packets :          81049992          0 pps
Drop bytes    :              0          0 bps
Drop packets  :              0          0 pps
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0,
L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0,
MTU errors: 0, Resource errors: 0
Ingress queues: 4 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

0 best-effort          81049992          81049992          0

1 expedited-fo              0              0          0

2 assured-forw           0              0          0

3 network-cont           0              0          0

Egress queues: 4 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

0 best-effort          0              0          0

1 expedited-fo          0              0          0

2 assured-forw          0              0          0

3 network-cont          0              0          0

Active alarms : None
Active defects : None
PCS statistics
Bit errors          Seconds
Errorred blocks          0
Errorred blocks          0

```

```

MAC statistics:
Total octets          6970332384
Total packets        81050506
Unicast packets      81050000
Broadcast packets     506
Multicast packets     0
CRC/Align errors     0
FIFO errors          0
MAC control frames    0
MAC pause frames      0
Oversized frames      0
Jabber frames         0
Fragment frames       0
VLAN tagged frames    0
Code violations       0

Filter statistics:
Input packet count    81050506
Input packet rejects   506
Input DA rejects      0
Input SA rejects      0
Output packet count    0
Output packet pad count 0
Output packet error count 0
CAM destination filters: 0, CAM source filters: 0

Packet Forwarding Engine configuration:
Destination slot: 5

CoS information:
Direction : Output
CoS transmit queue    Bandwidth      Buffer Priority Limit
                        %      bps      %      usec
0 best-effort         95    950000000  95      0      low  none
3 network-control      5     50000000   5      0      low  none

Direction : Input
CoS transmit queue    Bandwidth      Buffer Priority Limit
                        %      bps      %      usec
0 best-effort         95    950000000  95      0      low  none
3 network-control      5     50000000   5      0      low  none

Logical interface xe-5/0/0.0 (Index 71) (SNMP ifIndex 95) (Generation 195)
Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 ] Encapsulation: ENET2
Egress account overhead: 100
Ingress account overhead: 90

Traffic statistics:
Input bytes : 0
Output bytes : 46
Input packets: 0
Output packets: 1

IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

Local statistics:
Input bytes : 0
Output bytes : 46
Input packets: 0
Output packets: 1

Transit statistics:
Input bytes : 0
Output bytes : 0

```

```

Input packets:                0                0 pps
Output packets:               0                0 pps
IPv6 transit statistics:
  Input bytes :                0
  Output bytes :               0
  Input packets:              0
  Output packets:             0
Protocol inet, MTU: 1500, Generation: 253, Route table: 0
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 192.1.1/24, Local: 192.1.1.1, Broadcast: 192.1.1.255,
Generation: 265
Protocol multiservice, MTU: Unlimited, Generation: 254, Route table: 0
  Flags: None
  Policer: Input: __default_arp_policer__

```

show interfaces extensive (10-Gigabit Ethernet, WAN PHY Mode)

```

user@host> show interfaces xe-1/0/0 extensive
Physical interface: xe-1/0/0, Enabled, Physical link is Up
Interface index: 141, SNMP ifIndex: 34, Generation: 47
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Loopback: Disabled
WAN-PHY mode
Source filtering: Disabled, Flow control: Enabled
Device flags : Present Running
Interface flags: SNMP-Traps 16384
Link flags : None
CoS queues : 4 supported
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:05:85:a2:10:9d, Hardware address: 00:05:85:a2:10:9d
Last flapped : 2005-07-07 11:22:34 PDT (3d 12:28 ago)
Statistics last cleared: Never
Traffic statistics:
  Input bytes :                0                0 bps
  Output bytes :               0                0 bps
  Input packets:              0                0 pps
  Output packets:             0                0 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  HS Link CRC errors: 0, HS Link FIFO overflows: 0,
  Resource errors: 0
Output errors:
  Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0,
  Aged packets: 0, FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0,
  Resource errors: 0
Queue counters:
  Queued packets  Transmitted packets  Dropped packets
0 best-effort    0                0                0
1 expedited-fo  0                0                0
2 assured-forw   0                0                0
3 network-cont   0                0                0
Active alarms : LOL, LOS, LBL
Active defects: LOL, LOS, LBL, SEF, AIS-L, AIS-P
PCS statistics
  Seconds  Count
Bit errors 0 0
Errored blocks 0 0
MAC statistics:
  Receive  Transmit
Total octets 0 0
Total packets 0 0
Unicast packets 0 0
Broadcast packets 0 0
Multicast packets 0 0

```

```

CRC/Align errors                0          0
FIFO errors                     0          0
MAC control frames              0          0
MAC pause frames               0          0
Oversized frames               0
Jabber frames                  0
Fragment frames                0
VLAN tagged frames             0
Code violations                 0
Filter statistics:
  Input packet count            0
  Input packet rejects          0
  Input DA rejects              0
  Input SA rejects              0
  Output packet count           0
  Output packet pad count       0
  Output packet error count     0
CAM destination filters: 0, CAM source filters: 0
PMA PHY:
  Seconds      Count  State
  PLL lock     0      0 OK
  PHY light    63159  1 Light Missing
WIS section:
  BIP-B1        0      0
  SEF           434430  434438 Defect Active
  LOS           434430  1 Defect Active
  LOF           434430  1 Defect Active
  ES-S          434430
  SES-S         434430
  SEFS-S        434430
WIS line:
  BIP-B2        0      0
  REI-L         0      0
  RDI-L         0      0 OK
  AIS-L         434430  1 Defect Active
  BERR-SF       0      0 OK
  BERR-SD       0      0 OK
  ES-L          434430
  SES-L         434430
  UAS-L         434420
  ES-LFE        0
  SES-LFE       0
  UAS-LFE       0
WIS path:
  BIP-B3        0      0
  REI-P         0      0
  LOP-P         0      0 OK
  AIS-P         434430  1 Defect Active
  RDI-P         0      0 OK
  UNEQ-P        0      0 OK
  PLM-P         0      0 OK
  ES-P          434430
  SES-P         434430
  UAS-P         434420
  ES-PFE        0
  SES-PFE       0
  UAS-PFE       0
Received path trace:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Transmitted path trace: orissa so-1/0/0
6f 72 69 73 73 61 20 73 6f 2d 31 2f 30 2f 30 00 orissa so-1/0/0.
Packet Forwarding Engine configuration:

```

```

Destination slot: 1
CoS information:
  CoS transmit queue      Bandwidth      Buffer      Priority  Limit
                           %      bps      %      bytes
  0 best-effort           95      950000000  95        0      low      none
  3 network-control       5       50000000   5         0      low      none

```

show interfaces extensive (10-Gigabit Ethernet, DWDM OTN PIC)

```

user@host> show interfaces ge-7/0/0 extensive
Physical interface: ge-7/0/0, Enabled, Physical link is Down
Interface index: 143, SNMP ifIndex: 508, Generation: 208
Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
Flow control: Enabled
Device flags   : Present Running Down
Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
Link flags     : None
Wavelength     : 1550.12 nm, Frequency: 193.40 THz
CoS queues     : 8 supported, 8 maximum usable queues
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:05:85:70:2b:72, Hardware address: 00:05:85:70:2b:72
Last flapped   : 2011-04-20 15:48:54 PDT (18:39:49 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes   : 0          0 bps
Output bytes  : 0          0 bps
Input packets : 0          0 pps
Output packets: 0          0 pps
IPv6 transit statistics:
Input bytes   : 0
Output bytes  : 0
Input packets : 0
Output packets: 0
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
FIFO errors: 0, Resource errors: 0
Output errors:
Carrier transitions: 2, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 best-effort           0          0          0

  1 expedited-fo         0          0          0

  2 assured-forw         0          0          0

  3 network-cont
Queue number:      Mapped forwarding classes
  0                best-effort
  1                expedited-forwarding
  2                assured-forwarding
  3                network-control
Active alarms  : LINK
Active defects : LINK
MAC statistics:
Total octets      Receive      Transmit
Total packets     0          0

```



```

Unicast packets                0                0
Broadcast packets              0                0
Multicast packets              0                0
CRC/Align errors               0                0
FIFO errors                    0                0
MAC control frames             0                0
MAC pause frames               0                0
Oversized frames               0
Jabber frames                  0
Fragment frames                0
VLAN tagged frames             0
Code violations                 0
Total octets                   0                0
Total packets                  0                0
Unicast packets                0                0
Broadcast packets              0                0
Multicast packets              0                0
CRC/Align errors               0                0
FIFO errors                    0                0
MAC control frames             0                0
MAC pause frames               0                0
Oversized frames               0
Jabber frames                  0
Fragment frames                0
VLAN tagged frames             0
Code violations                 0
OTN alarms                     :   None
OTN defects                    :   None
OTN FEC Mode                   :  GFEC
OTN Rate                       :  Fixed Stuff Bytes 11.0957Gbps
OTN Line Loopback              :  Enabled
OTN FEC statistics :
  Corrected Errors              0
  Corrected Error Ratio (      0 sec average)  0e-0
OTN FEC alarms:                Seconds    Count  State
  FEC Degrade                   0          0  OK
  FEC Excessive                 0          0  OK
OTN OC:                        Seconds    Count  State
  LOS                           2          1  OK
  LOF                          67164       2  Defect Active
  LOM                          67164       71  Defect Active
  Wavelength Lock               0          0  OK
OTN OTU:
  AIS                           0          0  OK
  BDI                          65919      4814  Defect Active
  IAE                          67158       1  Defect Active
  TTIM                          7          1  OK
  SF                           67164       2  Defect Active
  SD                           67164       3  Defect Active
  TCA-ES                        0          0  OK
  TCA-SES                       0          0  OK
  TCA-UAS                       80         40  OK
  TCA-BBE                       0          0  OK
  BIP                           0          0  OK
  BBE                           0          0  OK
  ES                            0          0  OK
  SES                           0          0  OK
  UAS                           587         0  OK
Received DAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Received SAPI:

```

```

00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Transmitted DAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
Transmitted SAPI:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
OTN Received Overhead Bytes:
  APS/PCC0: 0x02, APS/PCC1: 0x42, APS/PCC2: 0xa2, APS/PCC3: 0x48
  Payload Type: 0x03
OTN Transmitted Overhead Bytes:
  APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
  Payload Type: 0x03
Filter statistics:
  Input packet count                0
  Input packet rejects              0
  Input DA rejects                  0
  Input SA rejects                  0
  Output packet count                0
  Output packet pad count            0
  Output packet error count          0
  CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 7
CoS information:
  Direction : Output
  CoS transmit queue      Bandwidth      Buffer Priority
Limit                    %      bps      %      usec
0 best-effort            95      9500000000    95      0      low
none
3 network-control        5       500000000     5      0      low
none
...

```

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode)

```

user@host> show interfaces xe-7/0/0 extensive
Physical interface: xe-7/0/0, Enabled, Physical link is Up
  Interface index: 173, SNMP ifIndex: 212, Generation: 174
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
  Unidirectional: Enabled,
  Loopback: None, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
...

```

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Transmit-Only)

```

user@host> show interfaces xe-7/0/0-tx extensive
Physical interface: xe-7/0/0-tx, Enabled, Physical link is Up
  Interface index: 176, SNMP ifIndex: 137, Generation: 177
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
  Unidirectional: Tx-Only
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:05:85:73:e4:83, Hardware address: 00:05:85:73:e4:83
  Last flapped   : 2007-06-01 09:08:19 PDT (3d 02:31 ago)
  Statistics last cleared: Never
Traffic statistics:
  Input bytes :                0                0 bps

```

```

Output bytes :      322891152287160      9627472888 bps
Input packets:              0              0 pps
Output packets:    328809727380      1225492 pps

...

Filter statistics:
  Output packet count      328810554250
  Output packet pad count      0
  Output packet error count    0
...

Logical interface xe-7/0/0-tx.0 (Index 73) (SNMP ifIndex 138) (Generation 139)

Flags: SNMP-Traps Encapsulation: ENET2
Egress account overhead: 100
Ingress account overhead: 90
Traffic statistics:
  Input bytes :              0
  Output bytes :    322891152287160
  Input packets:              0
  Output packets:    328809727380
IPv6 transit statistics:
  Input bytes :              0
  Output bytes :              0
  Input packets:              0
  Output packets:              0
Local statistics:
  Input bytes :              0
  Output bytes :              0
  Input packets:              0
  Output packets:              0
Transit statistics:
  Input bytes :              0              0 bps
  Output bytes :    322891152287160      9627472888 bps
  Input packets:              0              0 pps
  Output packets:    328809727380      1225492 pps
IPv6 transit statistics:
  Input bytes :              0
  Output bytes :              0
  Input packets:              0
  Output packets:              0
Protocol inet, MTU: 1500, Generation: 147, Route table: 0
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.11.12/24, Local: 10.11.12.13, Broadcast: 10.11.12.255,
Generation: 141
  Protocol multiservice, MTU: Unlimited, Generation: 148, Route table: 0
  Flags: None
  Policer: Input: __default_arp_policer__

```

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only)

```

user@host> show interfaces xe-7/0/0-rx extensive
Physical interface: xe-7/0/0-rx, Enabled, Physical link is Up
  Interface index: 174, SNMP ifIndex: 118, Generation: 175
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
Unidirectional: Rx-Only
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues

```

```
Hold-times      : Up 0 ms, Down 0 ms
Current address: 00:05:85:73:e4:83, Hardware address: 00:05:85:73:e4:83
Last flapped    : 2007-06-01 09:08:22 PDT (3d 02:31 ago)
Statistics last cleared: Never
Traffic statistics:
  Input bytes :      322857456303482      9627496104 bps
  Output bytes :              0              0 bps
  Input packets:      328775413751      1225495 pps
  Output packets:              0              0 pps

...

Filter statistics:
  Input packet count      328775015056
  Input packet rejects    1
  Input DA rejects        0

...

Logical interface xe-7/0/0-rx.0 (Index 72) (SNMP ifIndex 120) (Generation 138)

Flags: SNMP-Traps Encapsulation: ENET2
Traffic statistics:
  Input bytes :      322857456303482
  Output bytes :              0
  Input packets:      328775413751
  Output packets:              0
IPv6 transit statistics:
  Input bytes :              0
  Output bytes :              0
  Input packets:              0
  Output packets:              0
Local statistics:
  Input bytes :              0
  Output bytes :              0
  Input packets:              0
  Output packets:              0
Transit statistics:
  Input bytes :      322857456303482      9627496104 bps
  Output bytes :              0              0 bps
  Input packets:      328775413751      1225495 pps
  Output packets:              0              0 pps
IPv6 transit statistics:
  Input bytes :              0
  Output bytes :              0
  Input packets:              0
  Output packets:              0
Protocol inet, MTU: 1500, Generation: 145, Route table: 0
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 192.1.1/24, Local: 192.1.1.1, Broadcast: 192.1.1.255,
  Generation: 139
Protocol multiservice, MTU: Unlimited, Generation: 146, Route table: 0
  Flags: None
  Policer: Input: __default_arp_policer__
```

show interfaces (Fast Ethernet)

Syntax	<pre>show interfaces <i>interface-type</i> <brief detail extensive terse> <descriptions> <media> <snmp-index <i>snmp-index</i>> <statistics></pre>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display status information about the specified Fast Ethernet interface.
Options	<p><i>interface-type</i>—On M Series and T Series routers, the interface type is <i>fe-fpc/pic/port</i>. On the J Series routers, the interface type is <i>fe-pim/O/port</i>.</p> <p><i>brief detail extensive terse</i>—(Optional) Display the specified level of output.</p> <p><i>descriptions</i>—(Optional) Display interface description strings.</p> <p><i>media</i>—(Optional) Display media-specific information about network interfaces.</p> <p><i>snmp-index snmp-index</i>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><i>statistics</i>—(Optional) Display static interface statistics.</p>
Required Privilege Level	view
List of Sample Output	<p>show interfaces (Fast Ethernet) on page 232</p> <p>show interfaces brief (Fast Ethernet) on page 233</p> <p>show interfaces detail (Fast Ethernet) on page 233</p> <p>show interfaces extensive (Fast Ethernet) on page 233</p>
Output Fields	<p>Table 5 on page 219 lists the output fields for the show interfaces Fast Ethernet command. Output fields are listed in the approximate order in which they appear.</p>

Table 5: show interfaces Fast Ethernet Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none

Table 5: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Link-mode	Type of link connection configured for the physical interface: Full-duplex or Half-duplex	extensive
Speed	Speed at which the interface is running.	All levels
Loopback	Loopback status: Enabled or Disabled . If loopback is enabled, type of loopback: Local or Remote .	All levels
Source filtering	Source filtering status: Enabled or Disabled .	All levels
LAN-PHY mode	10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.	All levels
WAN-PHY mode	10-Gigabit Ethernet interface operating in Wide Area Network Physical Layer Device (WAN PHY) mode. WAN PHY allows 10-Gigabit Ethernet wide area links to use fiber-optic cables and other devices intended for SONET/SDH.	All levels
Unidirectional	Unidirectional link mode status for 10-Gigabit Ethernet interface: Enabled or Disabled for parent interface; Rx-only or Tx-only for child interfaces.	All levels
Flow control	Flow control status: Enabled or Disabled .	All levels
Auto-negotiation	(Gigabit Ethernet interfaces) Autonegotiation status: Enabled or Disabled .	All levels
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: <ul style="list-style-type: none"> • Online—Autonegotiation is manually configured as online. • Offline—Autonegotiation is manually configured as offline. 	All levels
Device flags	Information about the physical device. Possible values are described in the "Device Flags" section under <i>Common Output Fields Description</i> .	All levels
Interface flags	Information about the interface. Possible values are described in the "Interface Flags" section under <i>Common Output Fields Description</i> .	All levels
Link flags	Information about the link. Possible values are described in the "Links Flags" section under <i>Common Output Fields Description</i> .	All levels
Wavelength	(10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).	All levels

Table 5: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Frequency	(10-Gigabit Ethernet DWDM interfaces only) Displays the frequency associated with the configured wavelength, in terahertz (THz).	All levels
CoS queues	Number of CoS queues configured.	detail extensive none
Schedulers	(GigabitEthernet intelligent queuing 2 (IQ2) interfaces only) Number of CoS schedulers configured.	extensive
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	detail extensive
Current address	Configured MAC address.	detail extensive none
Hardware address	Hardware MAC address.	detail extensive none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .	detail extensive none
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output Rate	Output rate in bps and pps.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. <p>Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes.</p> <p>For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 31 under the show interfaces (10-Gigabit Ethernet) command.</p>	detail extensive

Table 5: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Input errors	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the ignore-l3-incompletes statement. • L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame. • L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • FIFO errors—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • Resource errors—Sum of transmit drops. 	extensive

Table 5: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Output errors	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Collisions—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • FIFO errors—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	extensive
Egress queues	Total number of egress queues supported on the specified interface.	detail extensive
Queue counters (Egress)	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	detail extensive
Ingress queues	Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.	extensive
Queue counters (Ingress)	<p>CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	extensive

Table 5: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Active alarms and Active defects	<p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the routing device configuration, an alarm can ring the red or yellow alarm bell on the routing device, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value None or Link.</p> <ul style="list-style-type: none"> • None—There are no active defects or alarms. • Link—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning. 	detail extensive none
OTN FEC statistics	<p>The forward error correction (FEC) counters provide the following statistics:</p> <ul style="list-style-type: none"> • Corrected Errors—The count of corrected errors in the last second. • Corrected Error Ratio—The corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits. 	
PCS statistics	<p>(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device.</p> <ul style="list-style-type: none"> • Bit errors—High bit error rate. Indicates the number of bit errors when the PCS receiver is operating in normal mode. • Errored blocks—Loss of block lock. The number of errored blocks when PCS receiver is operating in normal mode. 	detail extensive

Table 5: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
MAC statistics	<p>Receive and Transmit statistics reported by the PIC's MAC subsystem, including the following:</p> <ul style="list-style-type: none"> • Total octets and total packets—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. For more information, see Table 31 under the show interfaces (10-Gigabit Ethernet) command. • Unicast packets, Broadcast packets, and Multicast packets—Number of unicast, broadcast, and multicast packets. • CRC/Align errors—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • FIFO error—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning. • MAC control frames—Number of MAC control frames. • MAC pause frames—Number of MAC control frames with pause operational code. • Oversized frames—Number of frames that exceed 1518 octets. • Jabber frames—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • Fragment frames—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • VLAN tagged frames—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. • Code violations—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error." 	extensive
OTN Received Overhead Bytes	APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58 Payload Type: 0x08	extensive
OTN Transmitted Overhead Bytes	APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00 Payload Type: 0x08	extensive

Table 5: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Filter statistics	<p>Receive and Transmit statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected.</p> <ul style="list-style-type: none"> • Input packet count—Number of packets received from the MAC hardware that the filter processed. • Input packet rejects—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • Input DA rejects—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the routing device from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local routing device (which the routing device is rejecting). • Input SA rejects—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect. • Output packet count—Number of packets that the filter has given to the MAC hardware. • Output packet pad count—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured. • Output packet error count—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment. • CAM destination filters, CAM source filters—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields should be 0. 	extensive
PMA PHY	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • PHY Lock—Phase-locked loop • PHY Light—Loss of optical signal 	extensive

Table 5: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
WIS section	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • BIP-B1—Bit interleaved parity for SONET section overhead • SEF—Severely errored framing • LOL—Loss of light • LOF—Loss of frame • ES-S—Errored seconds (section) • SES-S—Severely errored seconds (section) • SEFS-S—Severely errored framing seconds (section) 	extensive
WIS line	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. State other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • BIP-B2—Bit interleaved parity for SONET line overhead • REI-L—Remote error indication (near-end line) • RDI-L—Remote defect indication (near-end line) • AIS-L—Alarm indication signal (near-end line) • BERR-SF—Bit error rate fault (signal failure) • BERR-SD—Bit error rate defect (signal degradation) • ES-L—Errored seconds (near-end line) • SES-L—Severely errored seconds (near-end line) • UAS-L—Unavailable seconds (near-end line) • ES-LFE—Errored seconds (far-end line) • SES-LFE—Severely errored seconds (far-end line) • UAS-LFE—Unavailable seconds (far-end line) 	extensive

Table 5: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
WIS path	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • BIP-B3—Bit interleaved parity for SONET section overhead • REI-P—Remote error indication • LOP-P—Loss of pointer (path) • AIS-P—Path alarm indication signal • RDI-P—Path remote defect indication • UNEQ-P—Path unequipped • PLM-P—Path payload (signal) label mismatch • ES-P—Errored seconds (near-end STS path) • SES-P—Severely errored seconds (near-end STS path) • UAS-P—Unavailable seconds (near-end STS path) • SES-PFE—Severely errored seconds (far-end STS path) • UAS-PFE—Unavailable seconds (far-end STS path) 	extensive

Table 5: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Autonegotiation information	<p>Information about link autonegotiation.</p> <ul style="list-style-type: none"> • Negotiation status: <ul style="list-style-type: none"> • Incomplete—Ethernet interface has the speed or link mode configured. • No autonegotiation—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation. • Complete—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner status—OK when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner: <ul style="list-style-type: none"> • Link mode—Depending on the capability of the attached Ethernet device, either Full-duplex or Half-duplex. • Flow control—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is None. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), and Symmetric/Asymmetric (link partner supports both PAUSE on receive and transmit or only PAUSE receive). • Remote fault—Remote fault information from the link partner—Failure indicates a receive link error. OK indicates that the link partner is receiving. Negotiation error indicates a negotiation error. Offline indicates that the link partner is going offline. • Local resolution—Information from the link partner: <ul style="list-style-type: none"> • Flow control—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), and Symmetric/Asymmetric (link partner supports both PAUSE on receive and transmit or only PAUSE receive). • Remote fault—Remote fault information. Link OK (no error detected on receive), Offline (local interface is offline), and Link Failure (link error detected on receive). 	extensive
Received path trace, Transmitted path trace	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other routing device manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the routing device at the other end of the fiber. The transmitted path trace value is the message that this routing device transmits.</p>	extensive
Packet Forwarding Engine configuration	<p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> • Destination slot—FPC slot number. 	extensive

Table 5: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
CoS information	Information about the CoS queue for the physical interface. <ul style="list-style-type: none"> • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Bandwidth bps—Bandwidth allocated to the queue (in bps). • Buffer %—Percentage of buffer space allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available. 	extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> .	All levels
VLAN-Tag	Rewrite profile applied to incoming or outgoing frames on the outer (Out) VLAN tag or for both the outer and inner (In) VLAN tags. <ul style="list-style-type: none"> • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • pop—The outer VLAN tag of the incoming frame is removed. • swap—The outer VLAN tag of the incoming frame is overwritten with the user specified VLAN tag information. • push-pop—An outer VLAN tag is pushed in front of the existing VLAN tag, and then removed. • push-push—Two VLAN tags are pushed in from the incoming frame. • swap-push—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame. • swap-swap—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user specified VLAN tag value. • pop-swap—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame. • pop-pop—Both the outer and inner VLAN tags of the incoming frame are removed. 	brief detail extensive none

Table 5: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Demux:	IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following: <ul style="list-style-type: none"> Source Family Inet Destination Family Inet 	detail extensive none
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family. Possible values are described in the “Protocol Field” section under <i>Common Output Fields Description</i> .	detail extensive none
MTU	Maximum transmission unit size on the logical interface.	detail extensive none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	detail extensive none
Traffic statistics	Number and rate of bytes and packets received and transmitted on the specified interface set. <ul style="list-style-type: none"> Input bytes, Output bytes—Number of bytes received and transmitted on the interface set Input packets, Output packets—Number of packets received and transmitted on the interface set. 	detail extensive
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	extensive
Local statistics	Number and rate of bytes and packets destined to the routing device.	extensive
Transit statistics	Number and rate of bytes and packets transiting the switch. <p>NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the Output bytes and Output packets interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.</p>	extensive
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive none
Flags	Information about protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	detail extensive
Donor interface	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	detail extensive none

Table 5: show interfaces Fast Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Preferred source address	(Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.	detail extensive none
Input Filters	Names of any input filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.	detail extensive
Output Filters	Names of any output filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.	detail extensive
Mac-Validate Failures	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	detail extensive none
Addresses, Flags	Information about the address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .	detail extensive none
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is inet , the IP address of the interface is also displayed.	brief
Flags	Information about address flag (possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i>).	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Sample Output

show interfaces (Fast Ethernet)

```

user@host> show interfaces fe-0/0/0
Physical interface: fe-0/0/0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 22
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Current address: 00:05:85:02:38:00, Hardware address: 00:05:85:02:38:00
  Last flapped   : 2006-01-20 14:50:58 PST (2w4d 00:44 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None
  Logical interface fe-0/0/0.0 (Index 66) (SNMP ifIndex 198)
    Flags: SNMP-Traps Encapsulation: ENET2

```

```

Protocol inet, MTU: 1500
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255

```

show interfaces brief (Fast Ethernet)

```

user@host> show interfaces fe-0/0/0 brief
Physical interface: fe-0/0/0, Enabled, Physical link is Up
Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
Logical interface fe-0/0/0.0
Flags: SNMP-Traps Encapsulation: ENET2
inet 10.10.10.1/24

```

show interfaces detail (Fast Ethernet)

```

user@host> show interfaces fe-0/0/0 detail
Physical interface: fe-0/0/0, Enabled, Physical link is Up
Interface index: 128, SNMP ifIndex: 22, Generation: 5391
Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
CoS queues : 4 supported, 4 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:05:85:02:38:00, Hardware address: 00:05:85:02:38:00
Last flapped : 2006-01-20 14:50:58 PST (2w4d 00:45 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes : 0 0 bps
Output bytes : 42 0 bps
Input packets: 0 0 pps
Output packets: 1 0 pps
Active alarms : None
Active defects : None
Logical interface fe-0/0/0.0 (Index 66) (SNMP ifIndex 198) (Generation 67)
Flags: SNMP-Traps Encapsulation: ENET2
Protocol inet, MTU: 1500, Generation: 105, Route table: 0
Flags: Is-Primary, Mac-Validate-Strict
Mac-Validate Failures: Packets: 0, Bytes: 0
Addresses, Flags: Is-Preferred Is-Primary
Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255,
Generation: 136

```

show interfaces extensive (Fast Ethernet)

```

user@host> show interfaces fe-0/0/0 extensive
Physical interface: fe-0/0/0, Enabled, Physical link is Up
Interface index: 128, SNMP ifIndex: 22, Generation: 5391
Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed:
100mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
CoS queues : 4 supported, 4 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:05:85:02:38:00, Hardware address: 00:05:85:02:38:00
Last flapped : 2006-01-20 14:50:58 PST (2w4d 00:46 ago)
Statistics last cleared: Never

```

```

Traffic statistics:
Input bytes :          0          0 bps
Output bytes :         42          0 bps
Input packets:         0          0 pps
Output packets:        1          0 pps
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  FIFO errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 3, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Active alarms : None
Active defects : None
MAC statistics:
Total octets          Receive      Transmit
Total packets         0            1
Unicast packets       0            0
Broadcast packets     0            1
Multicast packets     0            0
CRC/Align errors      0            0
FIFO errors           0            0
MAC control frames    0            0
MAC pause frames      0            0
Oversized frames      0
Jabber frames         0
Fragment frames       0
VLAN tagged frames    0
Code violations        0
Filter statistics:
Input packet count    0
Input packet rejects  0
Input DA rejects      0
Input SA rejects      0
Output packet count   1
Output packet pad count 0
Output packet error count 0
CAM destination filters: 1, CAM source filters: 0
Autonegotiation information:
Negotiation status: Complete
Link partner:
  Link partner: Full-duplex, Flow control: None, Remote fault: Ok
Local resolution:
Packet Forwarding Engine configuration:
Destination slot: 0
CoS information:
      Bandwidth      Buffer Priority  Limit
      %             bps  %         usec
0 best-effort      95   950000000  95         0    low  none
3 network-control  5    50000000   5         0    low  none
Logical interface fe-0/0/0.0 (Index 66) (SNMP ifIndex 198) (Generation 67)
Flags: SNMP-Traps Encapsulation: ENET2
Protocol inet, MTU: 1500, Generation: 105, Route table: 0
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
  Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255,
  Generation: 136

```

show interfaces (Gigabit Ethernet)

Syntax	<pre>show interfaces <i>ge-fpc/pic/port</i> <brief detail extensive terse> <descriptions> <media> <snmp-index <i>snmp-index</i>> <statistics></pre>
Release Information	Command introduced before Junos OS Release 7.4.
Description	(M Series, T Series, and MX Series routers and EX Series switches only) Display status information about the specified Gigabit Ethernet interface.
Options	<p><i>ge-fpc/pic/port</i>—Display standard information about the specified Gigabit Ethernet interface.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>descriptions—(Optional) Display interface description strings.</p> <p>media—(Optional) Display media-specific information about network interfaces.</p> <p>snmp-index <i>snmp-index</i>—(Optional) Display information for the specified SNMP index of the interface.</p> <p>statistics—(Optional) Display static interface statistics.</p>
Additional Information	In a logical system, this command displays information only about the logical interfaces and not about the physical interfaces.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> Verifying and Managing Agent Circuit Identifier-Based Dynamic VLAN Configuration
List of Sample Output	<p>show interfaces (Gigabit Ethernet) on page 250</p> <p>show interfaces (Gigabit Ethernet on MX Series Routers) on page 250</p> <p>show interfaces extensive (Gigabit Ethernet on MX Series Routers showing interface transmit statistics configuration) on page 251</p> <p>show interfaces brief (Gigabit Ethernet) on page 251</p> <p>show interfaces detail (Gigabit Ethernet) on page 252</p> <p>show interfaces extensive (Gigabit Ethernet IQ2) on page 253</p> <p>show interfaces (Gigabit Ethernet Unnumbered Interface) on page 256</p> <p>show interfaces (ACI Interface Set Configured) on page 256</p>
Output Fields	<p>Table 6 on page 236 describes the output fields for the show interfaces (Gigabit Ethernet) command. Output fields are listed in the approximate order in which they appear. For Gigabit Ethernet IQ and IQE PICs, the traffic and MAC statistics vary by interface type. For more information, see Table 7 on page 249.</p>

Table 6: show interfaces Gigabit Ethernet Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels
Loopback	Loopback status: Enabled or Disabled . If loopback is enabled, type of loopback: Local or Remote .	All levels
Source filtering	Source filtering status: Enabled or Disabled .	All levels
LAN-PHY mode	10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.	All levels
WAN-PHY mode	10-Gigabit Ethernet interface operating in Wide Area Network Physical Layer Device (WAN PHY) mode. WAN PHY allows 10-Gigabit Ethernet wide area links to use fiber-optic cables and other devices intended for SONET/SDH.	All levels
Unidirectional	Unidirectional link mode status for 10-Gigabit Ethernet interface: Enabled or Disabled for parent interface; Rx-only or Tx-only for child interfaces.	All levels
Flow control	Flow control status: Enabled or Disabled .	All levels
Auto-negotiation	(Gigabit Ethernet interfaces) Autonegotiation status: Enabled or Disabled .	All levels
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: <ul style="list-style-type: none"> Online—Autonegotiation is manually configured as online. Offline—Autonegotiation is manually configured as offline. 	All levels
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> .	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .	All levels

Table 6: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Link flags	Information about the link. Possible values are described in the “Links Flags” section under <i>Common Output Fields Description</i> .	All levels
Wavelength	(10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).	All levels
Frequency	(10-Gigabit Ethernet DWDM interfaces only) Displays the frequency associated with the configured wavelength, in terahertz (THz).	All levels
CoS queues	Number of CoS queues configured.	detail extensive none
Schedulers	(Gigabit Ethernet intelligent queuing 2 [IQ2] interfaces only) Number of CoS schedulers configured.	extensive
Hold-times	Current interface hold-time up and hold-time down, in milliseconds (ms).	detail extensive
Current address	Configured MAC address.	detail extensive none
Hardware address	Hardware MAC address.	detail extensive none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .	detail extensive none
Input Rate	Input rate in bits per second (bps) and packets per second (pps). The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	None
Output Rate	Output rate in bps and pps. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	None
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive
Egress account overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for egress traffic.	detail extensive
Ingress account overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for ingress traffic.	detail extensive

Table 6: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Output bytes—Number of bytes transmitted on the interface. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. <p>Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes.</p> <p>For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 31 under the show interfaces (10-Gigabit Ethernet) command.</p>	detail extensive
Input errors	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the ignore-l3-incompletes statement. • L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame. • L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • FIFO errors—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • Resource errors—Sum of transmit drops. 	extensive

Table 6: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Output errors	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. <p>NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the Drops field does not always use the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p> <ul style="list-style-type: none"> • Collisions—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • FIFO errors—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	extensive
Egress queues	Total number of egress queues supported on the specified interface.	detail extensive
Queue counters (Egress)	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. <p>NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the Dropped packets field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>	detail extensive
Ingress queues	Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.	extensive

Table 6: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Queue counters (Ingress)	CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces. <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	extensive
Active alarms and Active defects	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the router configuration, an alarm can ring the red or yellow alarm bell on the router, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value None or Link . <ul style="list-style-type: none"> • None—There are no active defects or alarms. • Link—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning. 	detail extensive none
Interface transmit statistics	(On MX Series devices) Status of the interface-transmit-statistics configuration: Enabled or Disabled. <ul style="list-style-type: none"> • Enabled—When the interface-transmit-statistics statement is included in the configuration. If this is configured, the interface statistics show the actual transmitted load on the interface. • Disabled—When the interface-transmit-statistics statement is not included in the configuration. If this is not configured, the interface statistics show the offered load on the interface. 	detail extensive
OTN FEC statistics	The forward error correction (FEC) counters provide the following statistics: <ul style="list-style-type: none"> • Corrected Errors—The count of corrected errors in the last second. • Corrected Error Ratio—The corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits. 	detail extensive
PCS statistics	(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device. <ul style="list-style-type: none"> • Bit errors—High bit error rate. Indicates the number of bit errors when the PCS receiver is operating in normal mode. • Errored blocks—Loss of block lock. The number of errored blocks when the PCS receiver is operating in normal mode. 	detail extensive

Table 6: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
MAC statistics	<p>Receive and Transmit statistics reported by the PIC's MAC subsystem, including the following:</p> <ul style="list-style-type: none"> • Total octets and total packets—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. For more information, see Table 31 under the show interfaces (10-Gigabit Ethernet) command. • Unicast packets, Broadcast packets, and Multicast packets—Number of unicast, broadcast, and multicast packets. • CRC/Align errors—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • FIFO error—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning. • MAC control frames—Number of MAC control frames. • MAC pause frames—Number of MAC control frames with pause operational code. • Oversized frames—There are two possible conditions regarding the number of oversized frames: <ul style="list-style-type: none"> • Packet length exceeds 1518 octets, or • Packet length exceeds MRU • Jabber frames—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • Fragment frames—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets) and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • VLAN tagged frames—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. <p>NOTE: The 20-port Gigabit Ethernet MIC (MIC-3D-20GE-SFP) does not have hardware counters for VLAN frames. Therefore, the VLAN tagged frames field displays 0 when the show interfaces command is executed on a 20-port Gigabit Ethernet MIC. In other words, the number of VLAN tagged frames cannot be determined for the 20-port Gigabit Ethernet MIC.</p> <ul style="list-style-type: none"> • Code violations—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error." 	extensive
OTN Received Overhead Bytes	APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58 Payload Type: 0x08	extensive
OTN Transmitted Overhead Bytes	APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00 Payload Type: 0x08	extensive

Table 6: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Filter statistics	<p>Receive and Transmit statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected.</p> <ul style="list-style-type: none"> • Input packet count—Number of packets received from the MAC hardware that the filter processed. • Input packet rejects—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • Input DA rejects—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the router from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local router (which the router is rejecting). • Input SA rejects—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect. • Output packet count—Number of packets that the filter has given to the MAC hardware. • Output packet pad count—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured. • Output packet error count—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment. • CAM destination filters, CAM source filters—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields should be 0. 	extensive
PMA PHY	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • PHY Lock—Phase-locked loop • PHY Light—Loss of optical signal 	extensive

Table 6: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
WIS section	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information:</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • BIP-B1—Bit interleaved parity for SONET section overhead • SEF—Severely errored framing • LOL—Loss of light • LOF—Loss of frame • ES-S—Errored seconds (section) • SES-S—Severely errored seconds (section) • SEFS-S—Severely errored framing seconds (section) 	extensive
WIS line	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information:</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • BIP-B2—Bit interleaved parity for SONET line overhead • REI-L—Remote error indication (near-end line) • RDI-L—Remote defect indication (near-end line) • AIS-L—Alarm indication signal (near-end line) • BERR-SF—Bit error rate fault (signal failure) • BERR-SD—Bit error rate defect (signal degradation) • ES-L—Errored seconds (near-end line) • SES-L—Severely errored seconds (near-end line) • UAS-L—Unavailable seconds (near-end line) • ES-LFE—Errored seconds (far-end line) • SES-LFE—Severely errored seconds (far-end line) • UAS-LFE—Unavailable seconds (far-end line) 	extensive

Table 6: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
WIS path	<p>(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information:</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. Any state other than OK indicates a problem. <p>Subfields are:</p> <ul style="list-style-type: none"> • BIP-B3—Bit interleaved parity for SONET section overhead • REI-P—Remote error indication • LOP-P—Loss of pointer (path) • AIS-P—Path alarm indication signal • RDI-P—Path remote defect indication • UNEQ-P—Path unequipped • PLM-P—Path payload (signal) label mismatch • ES-P—Errored seconds (near-end STS path) • SES-P—Severely errored seconds (near-end STS path) • UAS-P—Unavailable seconds (near-end STS path) • SES-PFE—Severely errored seconds (far-end STS path) • UAS-PFE—Unavailable seconds (far-end STS path) 	extensive

Table 6: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Autonegotiation information	<p>Information about link autonegotiation.</p> <ul style="list-style-type: none"> • Negotiation status: <ul style="list-style-type: none"> • Incomplete—Ethernet interface has the speed or link mode configured. • No autonegotiation—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation. • Complete—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner status—OK when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. • Link partner—Information from the remote Ethernet device: <ul style="list-style-type: none"> • Link mode—Depending on the capability of the link partner, either Full-duplex or Half-duplex. • Flow control—Types of flow control supported by the link partner. For Gigabit Ethernet interfaces, types are Symmetric (link partner supports PAUSE on receive and transmit), Asymmetric (link partner supports PAUSE on transmit), Symmetric/Asymmetric (link partner supports PAUSE on receive and transmit or only PAUSE on transmit), and None (link partner does not support flow control). • Remote fault—Remote fault information from the link partner—Failure indicates a receive link error. OK indicates that the link partner is receiving. Negotiation error indicates a negotiation error. Offline indicates that the link partner is going offline. • Local resolution—Information from the local Ethernet device: <ul style="list-style-type: none"> • Flow control—Types of flow control supported by the local device. For Gigabit Ethernet interfaces, advertised capabilities are Symmetric/Asymmetric (local device supports PAUSE on receive and transmit or only PAUSE on receive) and None (local device does not support flow control). Depending on the result of the negotiation with the link partner, local resolution flow control type will display Symmetric (local device supports PAUSE on receive and transmit), Asymmetric (local device supports PAUSE on receive), and None (local device does not support flow control). • Remote fault—Remote fault information. Link OK (no error detected on receive), Offline (local interface is offline), and Link Failure (link error detected on receive). 	extensive
Received path trace, Transmitted path trace	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other router manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the router at the other end of the fiber. The transmitted path trace value is the message that this router transmits.	extensive
Packet Forwarding Engine configuration	<p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> • Destination slot—FPC slot number. 	extensive

Table 6: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
CoS information	<p>Information about the CoS queue for the physical interface.</p> <ul style="list-style-type: none"> • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Bandwidth bps—Bandwidth allocated to the queue (in bps). • Buffer %—Percentage of buffer space allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available. 	extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under <i>Common Output Fields Description</i> .	All levels

Table 6: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
VLAN-Tag	<p>Rewrite profile applied to incoming or outgoing frames on the outer (Out) VLAN tag or for both the outer and inner (In) VLAN tags.</p> <ul style="list-style-type: none"> • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • pop—The outer VLAN tag of the incoming frame is removed. • swap—The outer VLAN tag of the incoming frame is overwritten with the user-specified VLAN tag information. • push—An outer VLAN tag is pushed in front of the existing VLAN tag. • push-push—Two VLAN tags are pushed in from the incoming frame. • swap-push—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame. • swap-swap—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user-specified VLAN tag value. • pop-swap—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame. • pop-pop—Both the outer and inner VLAN tags of the incoming frame are removed. 	brief detail extensive none
Demux	<p>IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following:</p> <ul style="list-style-type: none"> • Source Family Inet • Destination Family Inet 	detail extensive none
Encapsulation	Encapsulation on the logical interface.	All levels
ACI VLAN: Dynamic Profile	Name of the dynamic profile that defines the agent circuit identifier (ACI) interface set. If configured, the ACI interface set enables the underlying Ethernet interface to create dynamic VLAN subscriber interfaces based on ACI information.	brief detail extensive none
Protocol	Protocol family. Possible values are described in the “Protocol Field” section under <i>Common Output Fields Description</i> .	detail extensive none
MTU	Maximum transmission unit size on the logical interface.	detail extensive none
Dynamic Profile	(MX Series routers with Trio MPCs only) Name of the dynamic profile that was used to create this interface configured with a Point-to-Point Protocol over Ethernet (PPPoE) family.	detail extensive none
Service Name Table	(MX Series routers with Trio MPCs only) Name of the service name table for the interface configured with a PPPoE family.	detail extensive none
Max Sessions	(MX Series routers with Trio MPCs only) Maximum number of PPPoE logical interfaces that can be activated on the underlying interface.	detail extensive none

Table 6: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Duplicate Protection	(MX Series routers with Trio MPCs only) State of PPPoE duplicate protection: On or Off . When duplicate protection is configured for the underlying interface, a dynamic PPPoE logical interface cannot be activated when an existing active logical interface is present for the same PPPoE client.	detail extensive none
Direct Connect	State of the configuration to ignore DSL Forum VSAs: On or Off . When configured, the router ignores any of these VSAs received from a directly connected CPE device on the interface.	detail extensive none
AC Name	Name of the access concentrator.	detail extensive none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	detail extensive none
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the specified interface set.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Input packets, Output packets—Number of packets received and transmitted on the interface set. 	detail extensive
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	extensive
Local statistics	Number and rate of bytes and packets destined to the router.	extensive
Transit statistics	<p>Number and rate of bytes and packets transiting the switch.</p> <p>NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the Output bytes and Output packets interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.</p>	extensive
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive none
Flags	Information about protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	detail extensive
Donor interface	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	detail extensive none

Table 6: show interfaces Gigabit Ethernet Output Fields (*continued*)

Field Name	Field Description	Level of Output
Preferred source address	(Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.	detail extensive none
Input Filters	Names of any input filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parentheses next to all interfaces.	detail extensive
Output Filters	Names of any output filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parentheses next to all interfaces.	detail extensive
Mac-Validate Failures	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	detail extensive none
Addresses, Flags	Information about the address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .	detail extensive none
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is inet , the IP address of the interface is also displayed.	brief
Flags	Information about the address flag. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Table 7: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type

Interface Type	Sample Command	Byte and Octet Counts Include	Comments
Inbound physical interface	show interfaces ge-0/3/0 extensive	<p>Traffic statistics:</p> <p>Input bytes: 496 bytes per packet, representing the Layer 2 packet</p> <p>MAC statistics:</p> <p>Received octets: 500 bytes per packet, representing the Layer 2 packet + 4 bytes</p>	The additional 4 bytes are for the CRC.
Inbound logical interface	show interfaces ge-0/3/0.50 extensive	<p>Traffic statistics:</p> <p>Input bytes: 478 bytes per packet, representing the Layer 3 packet</p>	

Table 7: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type (*continued*)

Interface Type	Sample Command	Byte and Octet Counts Include	Comments
Outbound physical interface	show interfaces ge-0/0/0 extensive	Traffic statistics: Input bytes: 490 bytes per packet, representing the Layer 3 packet + 12 bytes MAC statistics: Received octets: 478 bytes per packet, representing the Layer 3 packet	For input bytes, the additional 12 bytes include 6 bytes for the destination MAC address plus 4 bytes for VLAN plus 2 bytes for the Ethernet type.
Outbound logical interface	show interfaces ge-0/0/0.50 extensive	Traffic statistics: Input bytes: 478 bytes per packet, representing the Layer 3 packet	

Sample Output

show interfaces (Gigabit Ethernet)

```

user@host> show interfaces ge-3/0/2
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Interface index: 167, SNMP ifIndex: 35
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Current address: 00:05:85:4a:e9:7c, Hardware address: 00:05:85:4a:e9:7c
  Last flapped   : 2006-08-10 17:25:10 PDT (00:01:08 ago)
  Input rate      : 0 bps (0 pps)
  Output rate     : 0 bps (0 pps)
  Ingress rate at Packet Forwarding Engine : 0 bps (0 pps)
  Ingress drop rate at Packet Forwarding Engine : 0 bps (0 pps)
  Active alarms   : None
  Active defects  : None

Logical interface ge-3/0/2.0 (Index 72) (SNMP ifIndex 69)
  Flags: SNMP-Traps 0x4000
  VLAN-Tag [ 0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push
  0x8100.512 0x8100.513)
  Encapsulation: VLAN-CCC
  Egress account overhead: 100
  Ingress account overhead: 90
  Input packets : 0
  Output packets: 0
  Protocol ccc, MTU: 1522
  Flags: Is-Primary

```

show interfaces (Gigabit Ethernet on MX Series Routers)

```

user@host> show interfaces ge-2/2/2
Physical interface: ge-2/2/2, Enabled, Physical link is Up
  Interface index: 156, SNMP ifIndex: 188
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, MAC-REWRITE Error: None,
  Loopback: Disabled,

```

```

Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
Link flags     : None
CoS queues    : 8 supported, 4 maximum usable queues
Schedulers    : 0
Current address: 00:1f:12:b7:d7:c0, Hardware address: 00:1f:12:b7:d6:76
Last flapped   : 2008-09-05 16:44:30 PDT (3d 01:04 ago)
Input rate     : 0 bps (0 pps)
Output rate    : 0 bps (0 pps)
Active alarms  : None
Active defects : None
Logical interface ge-2/2/2.0 (Index 82) (SNMP ifIndex 219)
  Flags: SNMP-Traps 0x20000000 Encapsulation: Ethernet-Bridge
  Egress account overhead: 100
  Ingress account overhead: 90
  Input packets : 0
  Output packets: 0
  Protocol aenet, AE bundle: ae0.0    Link Index: 4

```

show interfaces extensive (Gigabit Ethernet on MX Series Routers showing interface transmit statistics configuration)

```

user@host> show interfaces ge-2/1/2 extensive | match "output|interface"
Physical interface: ge-2/1/2, Enabled, Physical link is Up
Interface index: 151, SNMP ifIndex: 530, Generation: 154
Interface flags: SNMP-Traps Internal: 0x4000
Output bytes   :      240614363944      772721536 bps
Output packets:      3538446506      1420444 pps
Direction : Output
Interface transmit statistics: Enabled

Logical interface ge-2/1/2.0 (Index 331) (SNMP ifIndex 955) (Generation 146)
Output bytes   :      195560312716      522726272 bps
Output packets:      4251311146      1420451 pps

```

show interfaces brief (Gigabit Ethernet)

```

user@host> show interfaces ge-3/0/2 brief
Physical interface: ge-3/0/2, Enabled, Physical link is Up
Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags   : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
Link flags     : None

Logical interface ge-3/0/2.0
  Flags: SNMP-Traps 0x4000
  VLAN-Tag [ 0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push
0x8100.512 0x8100.513)
  Encapsulation: VLAN-CCC
  ccc

Logical interface ge-3/0/2.32767
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ] Encapsulation: ENET2

```

show interfaces detail (Gigabit Ethernet)

```

user@host> show interfaces ge-3/0/2 detail
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Interface index: 167, SNMP ifIndex: 35, Generation: 177
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 4 supported, 4 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:05:85:4a:e9:7c, Hardware address: 00:05:85:4a:e9:7c
  Last flapped   : 2006-08-09 17:17:00 PDT (01:31:33 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes   : 0 0 bps
    Output bytes  : 0 0 bps
    Input packets : 0 0 pps
    Output packets: 0 0 pps
  Ingress traffic statistics at Packet Forwarding Engine:
    Input bytes   : 0 0 bps
    Input packets : 0 0 pps
    Drop bytes    : 0 0 bps
    Drop packets  : 0 0 pps
  Ingress queues: 4 supported, 4 in use
  Queue counters:

```

	Queued packets	Transmitted packets	Dropped packets
0 best-effort	0	0	0
1 expedited-fo	0	0	0
2 assured-forw	0	0	0
3 network-cont	0	0	0

```

  Egress queues: 4 supported, 4 in use
  Queue counters:

```

	Queued packets	Transmitted packets	Dropped packets
0 best-effort	0	0	0
1 expedited-fo	0	0	0
2 assured-forw	0	0	0
3 network-cont	0	0	0

```

  Active alarms : None
  Active defects : None

  Logical interface ge-3/0/2.0 (Index 72) (SNMP ifIndex 69) (Generation 140)
    Flags: SNMP-Traps 0x4000
    VLAN-Tag [0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530)
  Out(swap-push 0x8100.512 0x8100.513)
    Encapsulation: VLAN-CCC
    Egress account overhead: 100
    Ingress account overhead: 90
    Traffic statistics:
      Input bytes   : 0
      Output bytes  : 0

```

```

Input packets:          0
Output packets:         0
Local statistics:
Input bytes :           0
Output bytes :          0
Input packets:          0
Output packets:         0
Transit statistics:
Input bytes :           0          0 bps
Output bytes :          0          0 bps
Input packets:          0          0 pps
Output packets:         0          0 pps
Protocol ccc, MTU: 1522, Generation: 149, Route table: 0
Flags: Is-Primary

```

```

Logical interface ge-3/0/2.32767 (Index 71) (SNMP ifIndex 70)
(Generation 139)
Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ] Encapsulation: ENET2
Traffic statistics:
Input bytes :           0
Output bytes :          0
Input packets:          0
Output packets:         0
Local statistics:
Input bytes :           0
Output bytes :          0
Input packets:          0
Output packets:         0
Transit statistics:
Input bytes :           0          0 bps
Output bytes :          0          0 bps
Input packets:          0          0 pps
Output packets:         0          0 pps

```

show interfaces extensive (Gigabit Ethernet IQ2)

```

user@host> show interfaces ge-7/1/3 extensive
Physical interface: ge-7/1/3, Enabled, Physical link is Up
Interface index: 170, SNMP ifIndex: 70, Generation: 171
Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x4004000
Link flags : None
CoS queues : 8 supported, 4 maximum usable queues
Schedulers : 256
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:14:f6:30:5e:74, Hardware address: 00:14:f6:30:5e:74
Last flapped : 2007-11-07 21:31:41 PST (02:03:33 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes :          38910844056          7952 bps
Output bytes :           7174605          8464 bps
Input packets:         418398473           11 pps
Output packets:          78903           12 pps
IPv6 transit statistics:
Input bytes :           0
Output bytes :           0
Input packets:          0
Output packets:          0

```

Ingress traffic statistics at Packet Forwarding Engine:

```

Input bytes :          38910799145          7952 bps
Input packets:         418397956           11 pps
Drop bytes :              0              0 bps
Drop packets:           0              0 pps

```

Input errors:

```

Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
FIFO errors: 0, Resource errors: 0

```

Output errors:

```

Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

```

```

FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0

```

```

Ingress queues: 4 supported, 4 in use

```

Queue counters:	Queued packets	Transmitted packets	Dropped packets
0 best-effort	418390823	418390823	0
1 expedited-fo	0	0	0
2 assured-forw	0	0	0
3 network-cont	7133	7133	0

```

Egress queues: 4 supported, 4 in use

```

Queue counters:	Queued packets	Transmitted packets	Dropped packets
0 best-effort	1031	1031	0
1 expedited-fo	0	0	0
2 assured-forw	0	0	0
3 network-cont	77872	77872	0

```

Active alarms : None

```

```

Active defects : None

```

MAC statistics:

	Receive	Transmit
Total octets	38910844056	7174605
Total packets	418398473	78903
Unicast packets	408021893366	1026
Broadcast packets	10	12
Multicast packets	418398217	77865
CRC/Align errors	0	0
FIFO errors	0	0
MAC control frames	0	0
MAC pause frames	0	0
Oversized frames	0	
Jabber frames	0	
Fragment frames	0	
VLAN tagged frames	0	
Code violations	0	OTN Received Overhead Bytes:
APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58		
Payload Type: 0x08		

```

OTN Transmitted Overhead Bytes:

```

```

APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
Payload Type: 0x08

```

Filter statistics:

Input packet count	418398473
Input packet rejects	479
Input DA rejects	479


```

Input SA rejects                                0
Output packet count                            78903
Output packet pad count                        0
Output packet error count                      0
CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
Negotiation status: Complete
Link partner:
  Link mode: Full-duplex, Flow control: Symmetric/Asymmetric,
  Remote fault: OK
Local resolution:
  Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:
Destination slot: 7
CoS information:
Direction : Output
CoS transmit queue      Bandwidth      Buffer      Priority      Limit
                        %      bps      %      usec
0 best-effort           95      950000000  95      0
low none
3 network-control       5      500000000   5      0
low none
Direction : Input
CoS transmit queue      Bandwidth      Buffer      Priority      Limit
                        %      bps      %      usec
0 best-effort           95      950000000  95      0
low none
3 network-control       5      500000000   5      0
low none

Logical interface ge-7/1/3.0 (Index 70) (SNMP ifIndex 85) (Generation 150)
Flags: SNMP-Traps Encapsulation: ENET2
Traffic statistics:
Input bytes :      812400
Output bytes :    1349206
Input packets:      9429
Output packets:    9449
IPv6 transit statistics:
Input bytes :      0
Output bytes :      0
Input packets:      0
Output packets:     0
Local statistics:
Input bytes :      812400
Output bytes :    1349206
Input packets:      9429
Output packets:    9449
Transit statistics:
Input bytes :      0      7440 bps
Output bytes :      0      7888 bps
Input packets:      0      10 pps
Output packets:      0      11 pps
IPv6 transit statistics:
Input bytes :      0
Output bytes :      0
Input packets:      0
Output packets:     0
Protocol inet, MTU: 1500, Generation: 169, Route table: 0
Flags: Is-Primary, Mac-Validate-Strict
Mac-Validate Failures: Packets: 0, Bytes: 0
Addresses, Flags: Is-Preferred Is-Primary

```

```
Input Filters: F1-ge-3/0/1.0-in, F3-ge-3/0/1.0-in
Output Filters: F2-ge-3/0/1.0-out (53)
Destination: 10.74.2/24, Local: 10.74.2.2, Broadcast: 10.74.2.255,
Generation: 196
Protocol multiservice, MTU: Unlimited, Generation: 170, Route table: 0
Flags: Is-Primary
Policer: Input: __default_arp_policer__
```

NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics displayed in the **show interfaces** command output might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the interface counters. For detailed information, see the description of the logical interface **Transit statistics** fields in [Table 6 on page 236](#).

show interfaces (Gigabit Ethernet Unnumbered Interface)

```
user@host> show interfaces ge-3/2/0
Physical interface: ge-3/2/0, Enabled, Physical link is Up
  Interface index: 148, SNMP ifIndex: 50
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 4 maximum usable queues
  Current address: 00:14:f6:11:26:f8, Hardware address: 00:14:f6:11:26:f8
  Last flapped   : 2006-10-27 04:42:23 PDT (08:01:52 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 624 bps (1 pps)
  Active alarms  : None
  Active defects : None

Logical interface ge-3/2/0.0 (Index 67) (SNMP ifIndex 85)
  Flags: SNMP-Traps Encapsulation: ENET2
  Input packets : 0
  Output packets: 6
  Protocol inet, MTU: 1500
  Flags: Unnumbered
  Donor interface: lo0.0 (Index 64)
  Preferred source address: 22.22.22.22
```

show interfaces (ACI Interface Set Configured)

```
user@host> show interfaces ge-1/0/0.4001
Logical interface ge-1/0/0.4001 (Index 340) (SNMP ifIndex 548)
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.4001 ] Encapsulation: PPP-over-

Ethernet
ACI VLAN:
  Dynamic Profile: aci-vlan-set-profile
  PPPoE:
    Dynamic Profile: aci-vlan-pppoe-profile,
    Service Name Table: None,
    Max Sessions: 32000, Max Sessions VSA Ignore: Off,
    Duplicate Protection: On, Short Cycle Protection: Off,
    Direct Connect: Off,
    AC Name: nbc
```

Input packets : 9
Output packets: 8
Protocol multiservice, MTU: Unlimited

CHAPTER 6

ANCP Operational Mode Commands

- clear ancp neighbor
- clear ancp statistics
- clear ancp subscriber
- request ancp oam interface
- request ancp oam neighbor
- show ancp cos
- show ancp neighbor
- show ancp statistics
- show ancp subscriber
- show ancp summary
- show ancp summary neighbor
- show ancp summary subscriber

clear ancp neighbor

Syntax	clear ancp neighbor <ip-address <i>ip-address</i>> <system-name <i>mac-address</i>>
Release Information	Command introduced in Junos OS Release 9.4.
Description	Clear the ANCP agent connection with all ANCP neighbors or with the specified ANCP neighbor. This command deletes information for subscribers associated with the neighbor, causing the adjusted traffic rates to revert to the configured rate for the subscriber interfaces. The neighbor remains configured (its administrative state is <i>enabled</i>) and can reestablish adjacencies.
Options	none —Clear all ANCP neighbors. ip-address <i>ip-address</i> —(Optional) Clear the ANCP neighbor specified by the IP address. system-name <i>mac-address</i> —(Optional) Clear the ANCP neighbor specified by the MAC address.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show ancp neighbor on page 271
List of Sample Output	clear ancp neighbor on page 260 show ancp neighbor on page 260
Output Fields	When you enter this command, you are provided no feedback on the status of your request. You can enter the show ancp neighbor command before and after clearing the ANCP neighbors to verify the clear operation.

Sample Output

clear ancp neighbor

```
user@host> clear ancp neighbor
```

show ancp neighbor

The following sample output displays the connections with ANCP neighbors before and after the **clear ancp neighbor** command was issued.

```
user@host> show ancp neighbor
```

IP Address	MAC Address	State	Subscriber Count	Capabilities
10.10.10.2	ba:ad:be:ef:10:10	Established	5	Topo
12.12.12.2	ba:ad:be:ef:10:12	Established	5	Topo
13.13.13.2	ba:ad:be:ef:10:13	Established	5	Topo

14.14.14.2	ba:ad:be:ef:10:14	Established	5	Topo
------------	-------------------	-------------	---	------

```
user@host> clear ancp neighbor ip-address 10.10.10.2
```

```
user@host> show ancp neighbor
```

IP Address	MAC Address	State	Subscriber Count	Capabilities
12.12.12.2	ba:ad:be:ef:10:12	Established	5	Topo
13.13.13.2	ba:ad:be:ef:10:13	Established	5	Topo
14.14.14.2	ba:ad:be:ef:10:14	Established	5	Topo

clear ancp statistics

Syntax	clear ancp statistics <ip-address <i>ip-address</i>> <system-name <i>mac-address</i>>
Release Information	Command introduced in Junos OS Release 13.3.
Description	Clear current statistics accumulated by the ANCP agent for all ANCP neighbors or the specified neighbor.
Options	none —Clear all ANCP statistics. ip-address <i>ip-address</i> —(Optional) Clear statistics for the ANCP neighbor specified by the IP address. system-name <i>mac-address</i> —(Optional) Clear statistics for the ANCP neighbor specified by the MAC address.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show ancp neighbor on page 271
List of Sample Output	clear ancp statistics on page 262 show ancp neighbor on page 262
Output Fields	When you enter this command, you are provided no feedback on the status of your request. You can enter the show ancp neighbor command before and after clearing the ANCP neighbor statistics to verify the clear operation.

Sample Output

clear ancp statistics

```
user@host> clear ancp statistics
```

show ancp neighbor

The following sample output displays statistics for an ANCP neighbor before and after the **clear ancp statistics** command was issued.

```
user@host> show ancp neighbor ip-address 192.168.10.1 detail
Neighbor Information
  IP Address           : 192.168.10.1
  System Name          : 00:00:64:1b:01:02
  Up Time              : 38
  TCP Port             : 64959
  State                : Established
  Subscriber Count     : 7
  Capabilities         : Topology Discovery
  System Instance      : 11
```



```

Peer Instance                : 1
Adjacency Timer (in 100ms)   : 50
Peer Adjacency Timer (in 100ms) : 100
Partition Type                : 0
Partition Flag                : 1
Partition Identifier          : 0
Dead Timer                    : 22
Received Syn Count            : 47
Received Synack Count         : 48
Received Rstack Count         : 2
Received Ack Count            : 12
Received Port Up Count        : 8
Received Port Down Count      : 2
Received Other Count          : 0
Sent Syn Count                : 48
Sent Synack Count             : 47
Sent Rstack Count             : 1
Sent Ack Count                : 12
Max Discovery Limit Exceed Count : 0

```

```
user@host> clear ancp statistics ip-address 192.168.10.1
```

```
user@host> show ancp neighbor ip-address 192.168.10.1 detail
```

Neighbor Information

```

IP Address                   : 192.168.10.1
System Name                   : 00:00:64:1b:01:02
Up Time                       : 38
TCP Port                      : 64959
State                         : Established
Subscriber Count              : 7
Capabilities                   : Topology Discovery
System Instance               : 11
Peer Instance                 : 1
Adjacency Timer (in 100ms)    : 50
Peer Adjacency Timer (in 100ms) : 100
Partition Type                : 0
Partition Flag                : 1
Partition Identifier          : 0
Dead Timer                    : 22
Received Syn Count            : 0
Received Synack Count         : 0
Received Rstack Count         : 0
Received Ack Count            : 0
Received Port Up Count        : 0
Received Port Down Count      : 0
Received Other Count          : 0
Sent Syn Count                : 0
Sent Synack Count             : 0
Sent Rstack Count             : 0
Sent Ack Count                : 0
Max Discovery Limit Exceed Count : 0

```

clear ancp subscriber

Syntax	<code>clear ancp subscriber</code> <code><identifier <i>identifier-string</i>></code> <code><ip-address <i>ip-address</i>></code> <code>system-name <i>mac-address</i>></code>
Release Information	Command introduced in Junos OS Release 11.4.
Description	Clear the ANCP agent connection with all ANCP subscribers or with the specified ANCP subscriber. This command deletes information for the subscribers, causing the adjusted traffic rate to revert to the configured rate for the subscriber interface, but otherwise has no affect on ANCP neighbors.
Options	none —Clear all ANCP subscribers. identifier <i>identifier-string</i> —(Optional) Clear the ANCP subscriber identified by the access loop ID. ip-address <i>ip-address</i> —(Optional) Clear all ANCP subscribers on the neighbor specified by the IP address. system-name <i>mac-address</i> —(Optional) Clear all ANCP subscribers on the neighbor specified by the MAC address.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show ancp subscriber on page 284
List of Sample Output	show ancp subscriber brief on page 264 clear ancp subscriber on page 265
Output Fields	When you enter this command, you are provided no feedback on the status of your request. You can enter the show ancp subscriber command before and after clearing the ANCP neighbors to verify the clear operation.

Sample Output

show ancp subscriber brief

```
user@host> show ancp subscriber brief
Loop Identifier      Type      Interface      Rate      Neighbor
                    Kbps
port-1-10            VDSL2     set-ge-10410   64        10.10.10.2
port-1-11            VDSL2     set-ge-10411   64        11.11.11.2
port-2-10            VDSL2     ge-1/0/4.12    64        10.12.12.2
port-2-10            VDSL2     ge-1/0/4.12    64        10.12.12.3
port-2-11            VDSL2     ge-1/0/4.13    64        10.13.13.2
user@host> clear ancp subscriber identifier port-2-10
```

```
user@host> show ancp subscriber brief
```

Loop Identifier	Type	Interface	Rate Kbps	Neighbor
port-1-10	VDSL2	set-ge-10410	64	10.10.10.2
port-1-11	VDSL2	set-ge-10411	64	11.11.11.2
port-2-11	VDSL2	ge-1/0/4.13	64	10.13.13.2

clear ancp subscriber

```
user@host> clear ancp subscriber
```

request ancp oam interface

Syntax	request ancp oam interface <i>(interface-name interface-set set-name)</i> <count <i>count</i> > <timeout <i>duration</i> >
Release Information	Command introduced in Junos OS Release 11.4.
Description	Trigger the access node to run a loopback test on the local loop between the access node and the customer premises equipment. You must specify either an ANCP interface or an ANCP interface set. The access node responds to the NAS with the results of the test.
Options	<p>interface-name—Name of the ANCP interface on whose local loop the loopback test is run.</p> <p>interface-set set-name—Name of the ANCP interface set on whose local loop the loopback test is run.</p> <p>count count—(Optional) Number of times a loopback message is sent on the local loop. Range: 1 through 32. Default: 1.</p> <p>timeout duration—(Optional) Period of time in seconds that the NAS waits for a response to the OAM request. Range: 0 through 255. Default: 5.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• <i>Triggering ANCP OAM to Test the Local Loop</i>
List of Sample Output	request ancp oam interface on page 266
Output Fields	When you enter this command, you are provided feedback on the status of your request, including the result of the test, the response code, and the response string returned with the OAM response in the event of failure, an error code is displayed.

Sample Output

request ancp oam interface

```
user@host> request ancp oam interface ge-1/0/4.12 count 5 timeout 40
request succeeded
0x503 : DSL line status showtime
DEFAULT RESPONSE
```

request ancp oam neighbor

Syntax	<code>request ancp oam neighbor</code> (<code>ip-address <i>ip-address</i> system-name <i>neighbor-name</i></code>) <code>subscriber <i>identifier-string</i></code> < <code>count <i>count</i></code> > < <code>timeout <i>duration</i></code> >
Release Information	Command introduced in Junos OS Release 11.4.
Description	Trigger the access node to run a loopback test on the local loop between the access node and the customer premises equipment. You must specify both the access node and the subscriber. The access node responds to the NAS with the results of the test.
Options	<p><code>ip-address <i>ip-address</i></code>—IP address that specifies the access node on whose local loop the loopback test is run.</p> <p><code>system-name <i>neighbor-name</i></code>—System name that specifies the access node on whose local loop the loopback test is run.</p> <p><code>subscriber <i>identifier-string</i></code>—Access identifier that specifies the subscriber on whose local loop the loopback test is run.</p> <p><code>count <i>count</i></code>—(Optional) Number of times a loopback message is sent on the local loop. Range: 1 through 32. Default: 1.</p> <p><code>timeout <i>duration</i></code>—(Optional) Period of time in seconds that the NAS waits for a response to the OAM request. Range: 0 through 255. Default: 5.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>Triggering ANCP OAM to Test the Local Loop</i>
List of Sample Output	request ancp oam subscriber on page 267
Output Fields	When you enter this command, you are provided feedback on the status of your request, including the result of the test, the response code, and the response string returned with the OAM response in the event of failure, an error code is displayed.

Sample Output

request ancp oam subscriber

```
user@host> request ancp oam neighbor 10.10.10.1 subscriber "dslam port-1-11"
request succeeded
0x503 : DSL line status showtime
DEFAULT RESPONSE
```

show ancp cos

Syntax	show ancp cos <i>identifier</i> <last-update> <pending-update>
Release Information	Command introduced in Junos OS Release 9.4.
Description	Display information about the CoS state for subscriber traffic.
Options	identifier <i>identifier</i> —(Optional) Display information about the local loops for the specified access identifier. last-update —(Optional) Display the most recently updated CoS information. pending-update —(Optional) Display the pending update of CoS information.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show ancp neighbor on page 271• show ancp statistics on page 279• show ancp subscriber on page 284
List of Sample Output	show ancp cos on page 269 show ancp cos last-update on page 270 show ancp cos pending-update on page 270
Output Fields	Table 8 on page 268 lists the output fields for the show ancp cos command. Output fields are listed in the approximate order in which they appear.

Table 8: show ancp cos Output Fields

Field Name	Field Description
QoS Adjust Flag	State of QoS adjust: TRUE (configured) or FALSE (not configured). The adjustment factor (percentage) is displayed for each line type: ADSL , ADSL2 , ADSL2+ , VDSL , VDSL2 , SDSL .
Keepalive Timer	Interval between the keepalive messages that the ANCP agent sends to CoS.
Cos State	State of the interaction between the ANCP agent and CoS: <ul style="list-style-type: none">• ANCPD_COS_CONNECT_NEEDED• ANCPD_COS_CONNECT_PENDING• ANCPD_COS_CONNECT_DONE• ANCPD_COS_SESSION_SENT• ANCPD_COS_WRITE_READY

Table 8: show ancp cos Output Fields (*continued*)

Field Name	Field Description
Connect Time	Time at which the ANCP agent connected to CoS; useful for debugging.
Session Time	Time at which the ANCP agent sent a session connect message to CoS; useful for debugging.
Routing Instance Time	Time at which the ANCP agent sent the routing instance to CoS; useful for debugging.
Keepalive Time	Time at which the last keepalive message was sent.
Update Time	Time at which the shaping rate was last updated.
Type	Subscriber access type: ifl indicates that a single VLAN carries subscriber traffic and iflset indicates that a set of VLANs carries subscriber traffic.
Name	System-wide name of the particular subscriber access.
Index	Access identifier.
Pending Update	Actual downstream data rate to be applied next to this local loop, in Kbps.
Last Update	Actual downstream data rate last applied to this local loop, in Kbps.

Sample Output

show ancp cos

```
user@host> show ancp cos
```

```

Qos Adjust Flag:      TRUE
ADSL adjustment factor: 95
ADSL2 adjustment factor: 90
ADSL2+ adjustment factor:85
VDSL adjustment factor: 80
VDSL2 adjustment factor: 75
SDSL adjustment factor: 70
Keepalive Timer:      45 secs
Cos State:             WRITE_READY
Connect Time:          Not Set
Session Time:          Not Set
Routing Instance Time: Not Set
Keepalive Time:        Not Set
Update Time:           Not Set

```

Type	Name	Index	Pending Update	Last Update
iflset	set-ge-10410	1	None	64 Kbps
iflset	set-ge-10411	2	None	64 Kbps
ifl	ge-1/0/4.2	71	None	64 Kbps
ifl	ge-1/0/4.3	72	None	64 Kbps

show ancp cos last-update

```
user@host> show ancp cos last-update
```

```
Qos Adjust Flag:      TRUE
ADSL adjustment factor: 95
ADSL2 adjustment factor: 90
ADSL2+ adjustment factor:85
VDSL adjustment factor: 80
VDSL2 adjustment factor: 75
SDSL adjustment factor: 70
Keepalive Timer:      45 secs
Cos State:            WRITE_READY
Connect Time:         Mon Nov 17 15:03:01 2008
Session Time:         Mon Nov 17 15:03:13 2008
Routing Instance Time: Mon Nov 17 15:03:14 2008
Keepalive Time:       Wed Nov 19 15:32:14 2008
Update Time:         Mon Nov 17 15:03:15 2008
```

Type	Name	Index	Pending Update	Last Update
iflset	iflset0	1	None	64 Kbps
iflset	iflset1	2	None	64 Kbps

show ancp cos pending-update

```
user@host> show ancp cos pending-update
```

```
Qos Adjust Flag:      TRUE
ADSL adjustment factor: 95
ADSL2 adjustment factor: 90
ADSL2+ adjustment factor:85
VDSL adjustment factor: 80
VDSL2 adjustment factor: 75
SDSL adjustment factor: 70
Keepalive Timer:      45 secs
Cos State:            WRITE_READY
Connect Time:         Not Set
Session Time:         Not Set
Routing Instance Time: Not Set
Keepalive Time:       Not Set
Update Time:         Not Set
Keepalive Timer:      45 secs
Connect Time:         Mon Nov 17 15:03:01 2008
Session Time:         Mon Nov 17 15:03:13 2008
Routing Instance Time: Mon Nov 17 15:03:14 2008
Keepalive Time:       Wed Nov 19 15:32:29 2008
Rate Update Time:     Mon Nov 17 15:03:15 2008
```


show ancp neighbor

Syntax	show ancp neighbor <brief detail> <ip-address ip-address> <system-name mac-address>
Release Information	Command introduced in Junos OS Release 9.4.
Description	Display information about all ANCP neighbors or the specified ANCP neighbor, regardless of operational state.
Options	brief detail —(Optional) Display the specified level of detail. ip-address ip-address —(Optional) IP address of the ANCP neighbor (access node). system-name mac-address —(Optional) MAC address of the ANCP neighbor (access node).
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show ancp cos on page 268• show ancp subscriber on page 284
List of Sample Output	show ancp neighbor on page 274 show ancp neighbor detail on page 274 show ancp neighbor ip-address on page 276 show ancp neighbor system-name on page 277
Output Fields	Table 9 on page 271 lists the output fields for the show ancp neighbor command. Output fields are listed in the approximate order in which they appear.

Table 9: show ancp neighbor Output Fields

Field Name	Field Description	Level of Output
Version	Version of the ANCP implementation: <ul style="list-style-type: none">• 0x31—General Switch Management Protocol (GSMP) version 3, sub-version 1; ANCP version before <i>RFC 6320, Protocol for Access Node Control Mechanism in Broadband Networks</i>.• 0x32—ANCP version 1, defined in <i>RFC 6320, Protocol for Access Node Control Mechanism in Broadband Networks</i>.	brief detail none
IP Address	IP address of the ANCP neighbor.	brief detail none
System Name	MAC address of the ANCP neighbor.	detail

Table 9: show ancp neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
State	Operational state of the ANCP adjacency: <ul style="list-style-type: none"> Configured—The neighbor has been configured, but has never been in the Established state. An asterisk (*) is prefixed to the neighbor entry for this state. Establishing—Adjacency negotiations are in progress for the neighbor. An asterisk (*) is prefixed to the neighbor entry for this state. This state is rarely seen because the adjacency is established so quickly. Established—Adjacency negotiations have succeeded for the neighbor and an ANCP session has been established. Not Estblshed—Not Established; adjacency negotiations are ready to begin. Indicates that this neighbor previously had been in the Established state; that is, it has lost a previously established adjacency. An asterisk (*) is prefixed to the neighbor entry for this state. 	All levels
Up Time	How long the adjacency has been up in one of the following formats: <ul style="list-style-type: none"> <i>nwndnh</i>—number of weeks, days, and hours <i>nd hh:mm:ss</i>—number of days, hours, minutes, and seconds 	brief detail none
Subscriber Count	Number of subscribers associated with the ANCP neighbor (access local loop).	brief none
Capabilities	Negotiated ANCP capability: <ul style="list-style-type: none"> Topo—Topology discovery. OAM—Performance of local Operations Administration Maintenance (OAM) procedures on an access loop controlled by the router. 	All levels
TCP Port	TCP port on which ANCP messages are exchanged.	detail
System Instance	Number identifying the ANCP link instance from the edge device's perspective.	detail
Peer Instance	Number identifying the ANCP instance from the access node's perspective. This number is unique and changes when the node or link comes back up after going down.	detail
Timer	Adjacency timer value advertised by the ANCP peer in 100 ms increments; the interval between ANCP ACK messages. This value remains constant for the duration of an ANCP session.	detail
Partition Type	Number that identifies whether partitions are used and how the ID is negotiated: <ul style="list-style-type: none"> 0—No partition. 1—Fixed partition requested. 2—Fixed partition assigned. 	detail
Partition Flag	Number that specifies the type of partition requested: 1 (new adjacency) or 2 (recovered adjacency).	detail
Partition Identifier	Number that associates the ANCP message with a specific partition.	detail

Table 9: show ancp neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
Partition Adjacencies	Number of adjacencies that share the partition.	detail
Dead Timer	Remaining period that the edge device waits for adjacency packets from a neighbor before declaring the neighbor to be down. The maximum dead time value is three times the configured adjacency timer value. This field displays the current value based on the time that the last adjacency packet was received.	detail
Received Syn Count	Number of synchronization messages received from neighbors to maintain adjacencies.	detail
Received Synack Count	Number of synchronization acknowledgment messages received from neighbors in response to the node's synchronization messages.	detail
Received Rstack Count	Number of messages received from neighbors indicating that the link to the neighbor needs to be reset.	detail
Received Ack Count	Number of acknowledgment messages periodically received from neighbors after an adjacency has been established.	detail
Received Port Up Count	Number of status messages received from neighbors indicating that a port has transitioned to the up state.	detail
Received Port Down Count	Number of status messages received from neighbors indicating that a port has transitioned to the down state.	detail
Received Generic Resp Count	Number of generic response messages received from neighbors.	detail
Received Adjacency Update Count	Number of adjacency update messages received from neighbors.	detail
Received OAM Count	Number of OAM responses received from neighbors in reply to request commands.	detail
Received Other Count	Number of all other ANCP message packets received from neighbors that do not fit into one of the other categories.	detail
Sent Syn Count	Number of synchronization messages sent to neighbors to maintain adjacencies.	detail
Sent Synack Count	Number of synchronization acknowledgment messages sent to neighbors in response to the their synchronization messages.	detail
Sent Rstack Count	Number of messages sent to neighbors indicating that the link to the neighbor needs to be reset.	detail
Sent Ack Count	Number of acknowledgment messages periodically sent to neighbors after an adjacency has been established.	detail

Table 9: show ancp neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
Sent Generic Resp Count	Number of generic response messages sent to neighbors.	detail
Sent OAM Count	Number of OAM request commands sent to neighbors.	detail
Max Discovery Limit Exceed Count	Number of times that the maximum number of discovery table entries accepted from the neighbor has been exceeded.	detail
Result Codes	<p>Number of generic response messages sent to neighbors that include each of the following result codes:</p> <ul style="list-style-type: none"> • Invalid Request Message Count—A properly formed request message violated the protocol because of timing (such as a race condition) or direction of transmission. • Specified Port(s) Down Count—One or more of the specified ports are down because of a state mismatch between the router and an ANCP control application. • Out of Resources Count—ANCP is out of resources, probably not related to the access lines. This result code is sent only by an access node. • Request Msg Not Implemented Count— • Malformed Msg Count—Message is malformed because it was corrupted in transit or there was an implementation error at either end of the connection. • TLV Missing Count—One or more mandatory TLVs was missing from a request. • Invalid TLV Contents Count—The contents of one or more TLVs in the request do not match its required specification. • Non-Existent Port(s) Count—One or more of the ports specified in a request do not exist, possibly because of a configuration mismatch between the access node and the router or AAA. 	detail

Sample Output

show ancp neighbor

```

user@host> show ancp neighbor
  Version IP Address      PartID  State      Time      Subscriber
Capabilities
  0x31    10.0.1.3             0       Established 11:24      2          Topo
  0x31    10.0.1.5             0       Not Estblshd 2:45       2          Topo
* 0x0    100.0.0.2            0       Establishing 0           0
* 0x0    192.0.1.0            0       Configured   0           0
* 0x0    192.0.22.1           0       Configured   0           0

```

show ancp neighbor detail

```

user@host> show ancp neighbor detail
Neighbor Information
  Version      : 0x31
  IP Address   : 192.85.1.5
  System Name  : 00:10:94:00:00:01
  Up Time      : 26
  TCP Port     : 32666

```

```

State : Established
Subscriber Count : 4
Capabilities : Topo
System Instance : 2
Peer Instance : 20
Adjacency Timer (in 100ms) : 100
Peer Adjacency Timer (in 100ms) : 100
Partition Type : 0
Partition Flag : 1
Partition Identifier : 0
Partition Adjacencies : 0
Dead Timer : 23
Received Syn Count : 1
Received Synack Count : 1
Received Rstack Count : 0
Received Ack Count : 4
Received Port Up Count : 10
Received Port Down Count : 0
Received Generic Resp Count : 0
Received Adjacency Update Count : 0
Received OAM Count : 0
Received Other Count : 0
Sent Syn Count : 1
Sent Synack Count : 2
Sent Rstack Count : 0
Sent Ack Count : 3
Sent Generic Resp Count : 0
Sent OAM Count : 0
Max Discovery Limit Exceed Count : 0
Result Codes:
Invalid Request Message Count : 0 Received Sent
Specified Port(s) Down Count : 0 0
Out of Resources Count : 0 0
Request Msg Not Implemented Count: 0 0
Malformed Msg Count : 0 0
TLV Missing Count : 0 0
Invalid TLV Contents Count : 0 0
Non-Existent Port(s) Count : 0 0

Version : 0x32
IP Address : 192.168.9.1
System Name : 00:00:64:1c:01:02
Up Time : 36
TCP Port : 61408
State : Not Established
Subscriber Count : 1
Capabilities : Topology Discovery
System Instance : 12
Peer Instance : 1
Adjacency Timer (in 100ms) : 50
Peer Adjacency Timer (in 100ms) : 100
Partition Type : 0
Partition Flag : 1
Partition Identifier : 0
Partition Adjacencies : 0
Dead Timer : 23
Received Syn Count : 24
Received Synack Count : 20
Received Rstack Count : 2
Received Ack Count : 9

```

```

Received Port Up Count      : 5
Received Port Down Count    : 0
Received Generic Resp Count : 0
Received Adjacency Update Count : 0
Received OAM Responses Count : 2
Received Other Count        : 0
Sent Syn Count              : 20
Sent Synack Count           : 24
Sent Rstack Count           : 1
Sent Generic Resp Count     : 0
Sent Ack Count              : 9
Sent OAM Requests Count     : 4
Max Discovery Limit Exceed Count : 0
Result Codes:
Invalid Request Message Count : 0
Specified Port(s) Down Count : 0
Out of Resources Count        : 0
Request Msg Not Implemented Count: 0
Malformed Msg Count          : 0
TLV Missing Count            : 0
Invalid TLV Contents Count    : 0
Non-Existent Port(s) Count    : 0

```

show ancp neighbor ip-address

```
user@host> show ancp neighbor ip-address 192.85.1.5
```

Neighbor Information

```

Version      : 0x32
IP Address   : 192.85.1.5
System Name  : ba:ad:be:ef:10:10
Up Time      : 26
TCP Port     : 32666
State        : Established
Subscriber Count : 4
Capabilities  : Topo
System Instance : 2
Peer Instance  : 20
Adjacency Timer (in 100ms) : 100
Peer Adjacency Timer (in 100ms) : 100
Partition Type : 0
Partition Flag : 1
Partition Identifier : 0
Partition Adjacencies : 0
Dead Timer     : 23
Received Syn Count : 1
Received Synack Count : 1
Received Rstack Count : 0
Received Ack Count : 4
Received Port Up Count : 10
Received Port Down Count : 0
Received Generic Resp Count : 0
Received Adjacency Update Count : 0
Received OAM Count : 0
Received Other Count : 0
Sent Syn Count : 1
Sent Synack Count : 2
Sent Rstack Count : 0
Sent Ack Count : 3
Sent Generic Resp Count : 0
Sent OAM Count : 0

```

Max Discovery Limit Exceed Count	: 0	
Result Codes:	Received	Sent
Invalid Request Message Count	: 0	0
Specified Port(s) Down Count	: 0	0
Out of Resources Count	: 0	0
Request Msg Not Implemented Count	: 0	0
Malformed Msg Count	: 0	0
TLV Missing Count	: 0	0
Invalid TLV Contents Count	: 0	0
Non-Existent Port(s) Count	: 0	0

show ancp neighbor system-name

user@host> show ancp neighbor ba:ad:be:ef:10:10 detail

Neighbor Information

Version	: 0x31	
IP Address	: 10.100.0.1	
System Name	: 00:00:64:1b:01:02	
Up Time	: 19	
TCP Port	: 1028	
State	: Established	
Subscriber Count	: 2	
Capabilities	: Topology Discovery, OAM	
System Instance	: 1	
Peer Instance	: 10	
Adjacency Timer (in 100ms)	: 100	
Peer Adjacency Timer (in 100ms)	: 250	
Partition Type	: 0	
Partition Flag	: 1	
Partition Identifier	: 0	
Partition Adjacencies	: 0	
Dead Timer	: 55	
Received Syn Count	: 1	
Received Synack Count	: 1	
Received Rstack Count	: 0	
Received Ack Count	: 1	
Received Port Up Count	: 34	
Received Port Down Count	: 0	
Received Generic Resp Count	: 0	
Received Adjacency Update Count	: 0	
Received OAM Responses Count	: 2	
Received Other Count	: 0	
Sent Syn Count	: 1	
Sent Synack Count	: 1	
Sent Rstack Count	: 0	
Sent Ack Count	: 3	
Sent Generic Resp Count	: 0	
Sent OAM Requests Count	: 4	
Max Discovery Limit Exceed Count	: 3	
Result Codes:	Received	Sent
Invalid Request Message Count	: 0	0
Specified Port(s) Down Count	: 0	0
Out of Resources Count	: 0	0
Request Msg Not Implemented Count	: 0	0
Malformed Msg Count	: 0	0
TLV Missing Count	: 0	0
Invalid TLV Contents Count	: 0	0
Non-Existent Port(s) Count	: 0	0

show ancp statistics

Syntax	show ancp statistics <ip-address <i>ip-address</i>> <system-name <i>mac-address</i>>
Release Information	Command introduced in Junos OS Release 13.3.
Description	Display statistics for all ANCP neighbors (access nodes) or the specified ANCP neighbor.
Options	<p>none—Display statistics for all ANCP neighbors, including global statistics not show for individual neighbors.</p> <p>ip-address <i>ip-address</i>—(Optional) Display statistics for only the neighbor with the specified IP address.</p> <p>system-name <i>mac-address</i>—(Optional) Display statistics for only the neighbor with the specified MAC address.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show ancp cos on page 268 • show ancp neighbor on page 271 • show ancp subscriber on page 284
List of Sample Output	show ancp statistics on page 281 show ancp statistics ip-address on page 282 show ancp statistics system-name on page 282
Output Fields	Table 10 on page 279 lists the output fields for the show ancp statistics command. Output fields are listed in the approximate order in which they appear.

Table 10: show ancp statistics Output Fields

Field Name	Field Description
Number of neighbors	Total count of ANCP neighbors.
Number of subscribers	Total count of ANCP subscribers.
Accept Count	Number of neighbor TCP/IP sessions accepted on listener socket.
Accept Fail Count	Number of neighbor TCP/IP sessions that failed due to one of the following causes: session already exists, maximum number of ANCP connections exceeded, creation of session or neighbor failed, or protocol start failed.

Table 10: show ancp statistics Output Fields (*continued*)

Field Name	Field Description
No Config Accept Deny Count	Number of neighbor TCP/IP sessions that failed because the neighbor was not configured.
Received Syn Count	Number of synchronization messages received from neighbors to maintain adjacencies.
Received Synack Count	Number of synchronization acknowledgment messages received from neighbors in response to the node's synchronization messages.
Received Rstack Count	Number of messages received from neighbors indicating that the link to the neighbor needs to be reset.
Received Ack Count	Number of acknowledgment messages periodically received from neighbors after an adjacency has been established.
Received Port Up Count	Number of status messages received from neighbors indicating that a port has transitioned to the up state.
Received Port Down Count	Number of status messages received from neighbors indicating that a port has transitioned to the down state.
Received Generic Resp Count	Number of generic response messages received from neighbors.
Received Adjacency Update Count	Number of adjacency update messages received from neighbors.
Received OAM Count	Number of OAM responses received from neighbors in reply to request commands.
Received Other Count	Number of all other ANCP message packets received from neighbors that do not fit into one of the other categories.
Sent Syn Count	Number of synchronization messages sent to neighbors to maintain adjacencies.
Sent Synack Count	Number of synchronization acknowledgment messages sent to neighbors in response to the their synchronization messages.
Sent Rstack Count	Number of messages sent to neighbors indicating that the link to the neighbor needs to be reset.
Sent Ack Count	Number of acknowledgment messages periodically sent to neighbors after an adjacency has been established.
Sent Generic Resp Count	Number of generic response messages sent to neighbors.
Sent OAM Count	Number of OAM request commands sent to neighbors.

Table 10: show ancp statistics Output Fields (*continued*)

Field Name	Field Description
Result Codes	<p>Number of generic response messages sent to neighbors that include each of the following result codes:</p> <ul style="list-style-type: none"> • Invalid Request Message Count—A properly formed request messages violated the protocol because of timing (such as a race condition) or direction of transmission. • Specified Port(s) Down Count—One or more of the specified ports are down because of a state mismatch between the router and an ANCP control application. • Out of Resources Count—the ANCP agent is out of resources, probably not related to the access lines. This result code is sent only by an access node. • Request Msg Not Implemented Count— • Malformed Msg Count—Message is malformed because it was corrupted in transit or there was an implementation error at either end of the connection. • TLV Missing Count—One or more mandatory TLVs was missing from a request. • Invalid TLV Contents Count—The contents of one or more TLVs in the request do not match its required specification. • Non-Existent Port(s) Count—One or more of the ports specified in a request do not exist, possibly because of a configuration mismatch between the access node and the router or AAA.

Sample Output

show ancp statistics

```

user@host> show ancp statistics
Statistics
  Number of neighbors           : 4
  Number of subscribers         : 6
  Accept Count                  : 0
  Accept Fail Count             : 0
  No Config Accept Deny Count  : 0
  Received Syn Count            : 2
  Received Synack Count         : 1
  Received Rstack Count         : 0
  Received Ack Count            : 8
  Received Port Up Count        : 7
  Received Port Down Count      : 0
  Received Generic Resp Count   : 0
  Received Adjacency Update Count : 0
  Received OAM Count            : 0
  Received Other Count          : 0
  Sent Syn Count                : 1
  Sent Synack Count             : 1
  Sent Rstack Count             : 0
  Sent Ack Count                : 17
  Sent Generic Resp Count       : 0
  Sent OAM Count                : 4
Result Codes:
  Invalid Request Message Count : 0
  Specified Port(s) Down Count  : 0
  Out of Resources Count        : 0
  Request Msg Not Implemented Count: 0
Received Sent

```

Malformed Msg Count	: 0	0
TLV Missing Count	: 0	0
Invalid TLV Contents Count	: 0	0
Non-Existent Port(s) Count	: 0	0

show ancp statistics ip-address

```
user@host> show ancp statistics ip-address 10.0.0.1
Statistics
  Received Syn Count           : 2
  Received Synack Count        : 1
  Received Rstack Count        : 0
  Received Ack Count           : 8
  Received Port Up Count       : 7
  Received Port Down Count     : 0
  Received Generic Resp Count  : 0
  Received Adjacency Update Count : 0
  Received OAM Count           : 0
  Received Other Count         : 0
  Sent Syn Count               : 1
  Sent Synack Count            : 1
  Sent Rstack Count            : 0
  Sent Ack Count               : 17
  Sent Generic Resp Count      : 0
  Sent OAM Count               : 4
Result Codes:
  Received Sent
  Invalid Request Message Count : 0 0
  Specified Port(s) Down Count  : 0 0
  Out of Resources Count         : 0 0
  Request Msg Not Implemented Count: 0 0
  Malformed Msg Count           : 0 0
  TLV Missing Count             : 0 0
  Invalid TLV Contents Count     : 0 0
  Non-Existent Port(s) Count    : 0 0
```

show ancp statistics system-name

```
user@host> show ancp statistics system-name 00:00:64:1b:01:02
Statistics
  Received Syn Count           : 2
  Received Synack Count        : 1
  Received Rstack Count        : 0
  Received Ack Count           : 8
  Received Port Up Count       : 7
  Received Port Down Count     : 0
  Received Generic Resp Count  : 0
  Received Adjacency Update Count : 0
  Received OAM Count           : 0
  Received Other Count         : 0
  Sent Syn Count               : 1
  Sent Synack Count            : 1
  Sent Rstack Count            : 0
  Sent Ack Count               : 17
  Sent Generic Resp Count      : 0
  Sent OAM Count               : 4
Result Codes:
  Received Sent
  Invalid Request Message Count : 0 0
  Specified Port(s) Down Count  : 0 0
  Out of Resources Count         : 0 0
```

Request Msg Not Implemented Count:	0	0
Malformed Msg Count	: 0	0
TLV Missing Count	: 0	0
Invalid TLV Contents Count	: 0	0
Non-Existent Port(s) Count	: 0	0

show ancp subscriber

Syntax	show ancp subscriber <brief detail> <identifier <i>identifier</i>> <neighbor <i>ip-address</i>>
Release Information	Command introduced in Junos OS Release 9.4.
Description	Display information about active subscribers regardless of the subscriber's operational state, for all subscribers (local access loops), the subscriber associated with the access line specified by an ACL, or the subscriber associated with the specified ANCP neighbor (access node).
Options	<p>brief detail—(Optional) Display the specified level of detail.</p> <p>identifier <i>identifier</i>—(Optional) Display information about the subscriber associated with the access line specified by the access identifier.</p> <p>neighbor <i>ip-address</i>—(Optional) Display information about the subscribers connected to the access node specified by the IP address.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear ancp subscriber on page 264 • show ancp cos on page 268 • show ancp neighbor on page 271
List of Sample Output	show ancp subscriber on page 287 show ancp subscriber brief on page 287 show ancp subscriber detail on page 287 show ancp subscriber identifier identifier-string detail on page 288
Output Fields	Table 11 on page 284 lists the output fields for the show ancp subscriber command. Output fields are listed in the approximate order in which they appear.

Table 11: show ancp subscriber Output Fields

Field Name	Field Description	Level of Output
Loop Identifier	<p>Access loop identifier as sent by the access node and configured to map the subscriber to an interface.</p> <p>An asterisk (*) indicates that the information might be stale due to receiving a Port Down message with a DSL Line State of Idle.</p> <p>Two asterisks (**) indicate that the neighbor associated with the subscriber has lost its adjacency. In this case, the DSL Line State might be Established.</p>	brief none
DSL Line State	State of the DSL line: Idle , Showtime , or Silent .	brief detail

Table 11: show ancp subscriber Output Fields (*continued*)

Field Name	Field Description	Level of Output
Type	Type of digital subscriber line employed by the access node: ADSL1 , ADSL2 , ADSL2+ , VDSL1 , VDSL2 , SDSL , or UNKNOWN .	brief none
Interface	Name of the interface set or logical interface.	brief detail none
Rate Kbps	Actual downstream data rate for this local loop.	brief none
Neighbor	IP address of ANCP neighbor (access node).	brief none
Access Loop Circuit Identifier	<p>Access loop circuit identifier as sent by the access node and configured to map the subscriber to an interface.</p> <p>An asterisk (*) indicates that the information might be stale due to receiving a Port Down message with a DSL Line State of Idle.</p> <p>Two asterisks (**) indicate that the neighbor associated with the subscriber has lost its adjacency. In this case, the DSL Line State might be Established.</p>	detail
Neighbor IP Address	IP address of the ANCP neighbor (access node).	detail
Aggregate Circuit Identifier Binary	Binary identifier for the VLAN circuit ID.	detail
DSL Type	Type of digital subscriber line employed by the access node: ADSL1 , ADSL2 , ADSL2+ , VDSL1 , VDSL2 , SDSL , or UNKNOWN .	detail
DSL Line Data Link	Data link protocol employed on the access loop: AAL5 or Ethernet .	detail
DSL Line Encapsulation	<p>Encapsulation type on the access loop, for Ethernet only:</p> <ul style="list-style-type: none"> 0—NA, type not conveyed 1—Untagged Ethernet 2—Single-tagged Ethernet 	detail
DSL Line Encapsulation Payload	<p>Payload carried across the access loop:</p> <ul style="list-style-type: none"> 0—NA, type not conveyed 1—PPPoA LLC 2—PPPoA null 3—IPoA LLC 4—IPoA null 5—Ethernet over AAL5 LLC with FCS 6—Ethernet over AAL5 LLC without FCS 7—Ethernet over AAL5 null with FCS 8—Ethernet over AAL5 null without FCS 	detail
Interface Type	Type of interface employed for subscriber traffic: ifl for a single VLAN or interface-set for a configured group of VLANs.	detail

Table 11: show ancp subscriber Output Fields (*continued*)

Field Name	Field Description	Level of Output
Actual Net Data Upstream	Actual upstream data rate for this local loop, in Kbps.	detail
Actual Net Data Downstream	Actual downstream data rate for this local loop, in Kbps.	detail
Minimum Net Data Upstream	Minimum upstream data rate desired by the operator for this local loop, in Kbps.	detail
Minimum Net Data Downstream	Minimum downstream data rate desired by the operator for this local loop, in Kbps.	detail
Maximum Net Data Upstream	Maximum upstream data rate desired by the operator for this local loop, in Kbps.	detail
Maximum Net Data Downstream	Maximum downstream data rate desired by the operator for this local loop, in Kbps.	detail
Attainable Net Data Upstream	Maximum attainable upstream data rate for this local loop, in Kbps.	detail
Attainable Net Data Downstream	Maximum attainable downstream data rate for this local loop, in Kbps.	detail
Minimum Low Power Data Downstream	Minimum downstream data rate desired by the operator for this local loop in low power state, in Kbps.	detail
Minimum Low Power Data Upstream	Minimum upstream data rate desired by the operator for this local loop in low power state, in Kbps.	detail
Maximum Interleave Delay Downstream	Maximum interleaving delay for downstream data, in milliseconds.	detail
Maximum Interleave Delay Upstream	Maximum interleaving delay for upstream data, in milliseconds.	detail
Actual Interleave Delay Downstream	Actual interleaving delay for downstream data, in milliseconds.	detail
Actual Interleave Delay Upstream	Actual interleaving delay for upstream data, in milliseconds.	detail

Sample Output

show ancp subscriber

```
user@host> show ancp subscriber
```

Loop Identifier	DSL Line State	Type	Interface	Rate Kbps	Neighbor
**circuit 101	Idle	ADSL1	----	32	10.0.1.3
**circuit 102	Idle	ADSL1	----	32	10.0.1.3
circuit 301	Showtime	ADSL1	----	32	10.0.1.5
circuit 302	Showtime	ADSL1	----	32	10.0.1.5

show ancp subscriber brief

```
user@host> show ancp subscriber brief
```

Loop Identifier	Type	Interface	Rate Kbps	Neighbor
port-1-10	VDSL2	set-ge-10410	64	10.10.10.2
port-1-11	VDSL2	set-ge-10411	64	11.11.11.2
port-2-10	VDSL2	ge-1/0/4.12	64	10.12.12.2
port-2-11	VDSL2	ge-1/0/4.13	64	10.13.13.2

show ancp subscriber detail

```
user@host> show ancp subscriber detail
```

Subscriber Information

- * Access Loop Circuit Identifier : circuit 101
 - Neighbor IP Address : 10.0.1.3
 - Aggregate Circuit Identifier Binary : 0/0
 - DSL Type : ADSL1
 - DSL Line State : Idle
 - DSL Line Data Link : Data link 2
 - DSL Line Encapsulation : N/A
 - DSL Line Encapsulation Payload : N/A
 - Interface Type : N/A
 - Interface : ----
 - Actual Net Data Upstream : 32
 - Actual Net Data Downstream : 32
 - Minimum Net Data Upstream : 0
 - Minimum Net Data Downstream : 0
 - Maximum Net Data Upstream : 0
 - Maximum Net Data Downstream : 0
 - Attainable Net Data Upstream : 1024
 - Attainable Net Data Downstream : 8192
 - Minimum Low Power Data Downstream : 32
 - Minimum Low Power Data Upstream : 32
 - Maximum Interleave Delay Downstream : 20
 - Maximum Interleave Delay Upstream : 20
 - Actual Interleave Delay Downstream : 20
 - Actual Interleave Delay Upstream : 20
- * Access Loop Circuit Identifier: circuit 102
 - Neighbor IP Address : 10.0.1.3
 - Aggregate Circuit Identifier Binary : 0/0
 - DSL Type : ADSL1
 - DSL Line State : Idle
 - DSL Line Data Link : Data link 2
 - DSL Line Encapsulation : N/A
 - DSL Line Encapsulation Payload : N/A

```
Interface Type           : N/A
Interface                 : ----
Actual Net Data Upstream  : 32
Actual Net Data Downstream : 32
Minimum Net Data Upstream : 0
Minimum Net Data Downstream : 0
Maximum Net Data Upstream : 0
Maximum Net Data Downstream : 0
Attainable Net Data Upstream : 1024
Attainable Net Data Downstream : 8192
Minimum Low Power Data Downstream : 32
Minimum Low Power Data Upstream : 32
Maximum Interleave Delay Downstream : 20
Maximum Interleave Delay Upstream : 20
Actual Interleave Delay Downstream : 20
Actual Interleave Delay Upstream : 20
...
```

show ancp subscriber identifier identifier-string detail

```
user@host> show ancp subscriber identifier port-1-11 detail
```

```
Access Loop Identifier : port-1-11
Neighbor IP Address    : 10.11.11.2
Aggregate Circuit Identifier Binary : 0/0
DSL Type               : DSL 0
Interface Type         : interface-set
Interface              : set-ge-10411
DSL Line State         : Show Time
Actual Net Data Upstream : 64
Actual Net Data Downstream : 64
DSL Line Data Link      : AAL5
DSL Line Encapsulation : N/A
DSL Line Encapsulation Payload : N/A
Minimum Net Data Upstream : 64
Minimum Net Data Downstream : 64
Maximum Net Data Upstream : 64
Maximum Net Data Downstream : 64
Attainable Net Data Upstream : 64
Attainable Net Data Downstream : 64
Minimum Low Power Data Downstream : 64
Minimum Low Power Data Upstream : 64
Maximum Interleave Delay Downstream : 50
Maximum Interleave Delay Upstream : 50
Actual Interleave Delay Downstream : 50
Actual Interleave Delay Upstream : 50
```

show ancp summary

Syntax	show ancp summary
Release Information	Command introduced in Junos OS Release 13.1.
Description	Display a summary of the counts and states for all ANCP neighbors and subscribers.
Options	This command has no options.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show ancp neighbor on page 271 • show ancp summary neighbor on page 291 • show ancp subscriber on page 284 • show ancp summary subscriber on page 293
List of Sample Output	show ancp summary on page 290
Output Fields	Table 12 on page 289 lists the output fields for the show ancp summary command. Output fields are listed in the approximate order in which they appear.

Table 12: show ancp summary Output Fields

Field Name	Field Description
Configured	Number of ANCP neighbors in the Configured state; that is, that have been configured but never established.
Establishing	Number of ANCP neighbors in the Establishing state; that is, where negotiations are in progress.
Established	Number of ANCP neighbors in the Established state; that is, where negotiations have succeeded and the ANCP session has been established.
Not Estblshd	Number of ANCP neighbors in the Not Estblshd state; that is, that have lost a previously established adjacency and are ready to begin negotiations.
Total	Total number of ANCP neighbors; sum of neighbors in the Configured , Establishing , Established , and Not Estblshd states.
Showtime	Number of DSL lines in Showtime state.
Idle	Number of DSL lines in Idle state.
Silent	Number of DSL lines in Silent state.
Unknown	Number of DSL lines where the state is not Showtime , Idle , or Silent .

Table 12: show ancp summary Output Fields (*continued*)

Field Name	Field Description
Total	Total number of DSL lines (ANCP subscribers); sum of DSL lines in the Showtime , Idle , Silent , and Unknown states.

Sample Output

show ancp summary

```
user@host> show ancp summary
```

Neighbors Summary:

Configured	Establishing	Established	Not Established	Total
-----	-----	-----	-----	-----
22	0	2	0	24

Subscribers Summary:

Showtime	Idle	Silent	Unknown	Total
-----	-----	-----	-----	-----
4	0	0	0	4

show ancp summary neighbor

Syntax	show ancp summary neighbor <ip-address <i>ip-address</i> system-name <i>mac-address</i> >
Release Information	Command introduced in Junos OS Release 13.1.
Description	Display a summary of the counts and states for all ANCP neighbors and of the neighbor's subscribers when you specify a particular neighbor.
Options	ip-address <i>ip-address</i> —(Optional) IP address of the ANCP neighbor (access node). system-name <i>mac-address</i> —(Optional) MAC address of the ANCP neighbor (access node).
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show ancp summary on page 289 • show ancp subscriber on page 284 • show ancp summary subscriber on page 293
List of Sample Output	show ancp summary neighbor on page 292 show ancp summary neighbor (IP Address) on page 292 show ancp summary neighbor (MAC Address) on page 292
Output Fields	Table 13 on page 291 lists the output fields for the show ancp summary command. Output fields are listed in the approximate order in which they appear.

Table 13: show ancp summary neighbor Output Fields

Field Name	Field Description
Configured	Number of ANCP neighbors in the Configured state; that is, that have been configured but never established.
Establishing	Number of ANCP neighbors in the Establishing state; that is, where negotiations are in progress.
Established	Number of ANCP neighbors in the Established state; that is, where negotiations have succeeded and the ANCP session has been established.
Not Estblshd	Number of ANCP neighbors in the Not Estblshd state; that is, that have lost a previously established adjacency and are ready to begin negotiations.
Total	Total number of ANCP neighbors; sum of neighbors in the Configured , Establishing , Established , and Not Estblshd states.
Showtime	Number of DSL lines for the neighbor in Showtime state.

Table 13: show ancp summary neighbor Output Fields (*continued*)

Field Name	Field Description
Idle	Number of DSL lines for the neighbor in Idle state.
Silent	Number of DSL lines for the neighbor in Silent state.
Unknown	Number of DSL lines for the neighbor where the state is not Showtime , Idle , or Silent .
Total	Total number of DSL lines (ANCP subscribers); sum of DSL lines in the Showtime , Idle , Silent , and Unknown states.

Sample Output

show ancp summary neighbor

```
user@host> show ancp summary neighbor
```

Neighbors Summary:

Configured	Establishing	Established	Not Established	Total
22	0	2	0	24

show ancp summary neighbor (IP Address)

```
user@host> show ancp summary neighbor ip-address 192.168.10.1
```

Neighbor Summary:192.168.10.1 status Established

Subscribers Summary:

Show Time	Idle	Silent	Unknown	Total
6	0	0	0	6

show ancp summary neighbor (MAC Address)

```
user@host> show ancp summary neighbor system-name 00:00:64:1b:01:02
```

Neighbor Summary:00:00:64:1b:01:02 status Established

Subscribers Summary:

Show Time	Idle	Silent	Unknown	Total
5	1	2	0	8

show ancp summary subscriber

Syntax	show ancp summary subscriber
Release Information	Command introduced in Junos OS Release 13.1.
Description	Display a summary of the counts and states for all ANCP subscribers.
Options	This command has no options.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show ancp summary on page 289 • show ancp neighbor on page 271 • show ancp summary neighbor on page 291
List of Sample Output	show ancp summary subscriber on page 293
Output Fields	Table 14 on page 293 lists the output fields for the show ancp summary subscriber command. Output fields are listed in the approximate order in which they appear.

Table 14: show ancp summary subscriber Output Fields

Field Name	Field Description
Showtime	Number of DSL lines in Showtime state.
Idle	Number of DSL lines in Idle state.
Silent	Number of DSL lines in Silent state.
Unknown	Number of DSL lines where the state is not Showtime , Idle , or Silent .
Total	Total number of DSL lines (ANCP subscribers); sum of DSL lines in the Showtime , Idle , Silent , and Unknown states.

Sample Output

show ancp summary subscriber

```
user@host> show ancp summary subscriber
```

```
Subscribers Summary:
Show Time      Idle      Silent      Unknown      Total
-----
           8           1           0           1          10
```


CHAPTER 7

BFD Operational Mode Commands

- `clear bfd adaptation`
- `clear bfd session`
- `show bfd session`

clear bfd adaptation

Syntax	<code>clear bfd adaptation</code> <code><address session-address></code> <code><discriminator discr-number></code>
Release Information	Command introduced before Junos OS Release 7.4.
Description	<p>Clear adaptation for Bidirectional Forwarding Detection (BFD) sessions. BFD is a simple hello mechanism that detects failures in a network. Configured BFD interval timers can change, adapting to network situations. Use this command to return BFD interval timers to their configured values.</p> <p>The clear bfd adaptation command is hitless, meaning that the command does not affect traffic flow on the routing device.</p>
Options	<p>none—Clear adaptation for all BFD sessions.</p> <p>address session-address—(Optional) Clear adaptation for all BFD sessions matching the specified address.</p> <p>discriminator discr-number—(Optional) Clear adaptation for the local BFD session matching the specified discriminator.</p>
Additional Information	For more information, see the description of the bfd-liveness-detection configuration statement in the <i>Junos Routing Protocols Configuration Guide</i> .
Required Privilege Level	clear
List of Sample Output	clear bfd adaptation on page 296
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear bfd adaptation

```
user@host> clear bfd adaptation
```

clear bfd session

Syntax	clear bfd session <address <i>session-address</i> > <discriminator <i>discr-number</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch and QFX Series)	clear bfd session <address <i>session-address</i> > <discriminator <i>discr-number</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Drop one or more Bidirectional Forwarding Detection (BFD) sessions.
Options	<p>none—Drop all BFD sessions.</p> <p>address <i>session-address</i>—(Optional) Drop all BFD sessions matching the specified address.</p> <p>discriminator <i>discr-number</i>—(Optional) Drop the local BFD session matching the specified discriminator.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show bfd session on page 298
List of Sample Output	clear bfd session on page 297
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear bfd session

```
user@host> clear bfd session
```

show bfd session

Syntax	<pre>show bfd session <brief detail extensive summary> <address address> <client rsvp-oam (brief detail extensive summary) vpls-oam (brief detail extensive instance instance-name summary)> <discriminator discriminator> <logical-system (all logical-system-name)> <prefix address></pre>
Syntax (EX Series Switch and QFX Series)	<pre>show bfd session <brief detail extensive summary> <address address> <client rsvp-oam (brief detail extensive summary) vpls-oam (brief detail extensive instance instance-name summary)> <discriminator discriminator> <prefix address></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Options discriminator and address introduced in Junos OS Release 8.2.</p> <p>Option prefix introduced in Junos OS Release 9.0.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p> <p>Option client introduced in Junos OS Release 12.3R3.</p>
Description	Display information about active Bidirectional Forwarding Detection (BFD) sessions.
Options	<p>none—(Same as brief) Display information about active BFD sessions.</p> <p>brief detail extensive summary—(Optional) Display the specified level of output.</p> <p>address address—(Optional) Display information about the BFD session for the specified neighbor address.</p> <p>client rsvp-oam (brief detail extensive summary) vpls-oam (brief detail extensive instance instance-name summary)—(Optional) Display information about RSVP-OAM or VPLS-OAM BFD sessions in the specified level of output. For VPLS-OAM, display the specified level of output or display information about all of the BFD sessions for the specified VPLS routing instance.</p> <p>discriminator discriminator—(Optional) Display information about the BFD session using the specified local discriminator.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>prefix address—(Optional) Display information about all of the BFD sessions for the specified LDP forwarding equivalence class (FEC).</p>

Required Privilege Level view

Related Documentation

- [clear bfd session on page 297](#)
- *Examples: Configuring BFD for Static Routes*
- *Example: Configuring BFD for OSPF*
- *Example: Configuring BFD for BGP*
- *Configuring PIM and the Bidirectional Forwarding Detection (BFD) Protocol*
- *Example: Configuring BFD for IS-IS*

List of Sample Output

- [show bfd session on page 302](#)
- [show bfd session brief on page 303](#)
- [show bfd session detail on page 303](#)
- [show bfd session detail \(with Authentication\) on page 303](#)
- [show bfd session address extensive on page 303](#)
- [show bfd session client rsvp-oam on page 304](#)
- [show bfd session client vpls-oam summary on page 304](#)
- [show bfd session client vpls-oam instance instance-name on page 304](#)
- [show bfd session extensive on page 304](#)
- [show bfd session extensive \(with Authentication\) on page 305](#)
- [show bfd session summary on page 305](#)

Output Fields [Table 15 on page 299](#) describes the output fields for the **show bfd session** command. Output fields are listed in the approximate order in which they appear.

Table 15: show bfd session Output Fields

Field Name	Field Description	Level of Output
Address	Address on which the BFD session is active.	brief detail extensive none
State	State of the BFD session: Up , Down , Init (initializing), or Failing .	brief detail extensive none
Interface	Interface on which the BFD session is active.	brief detail extensive none
Detect Time	Negotiated time interval, in seconds, used to detect BFD control packets.	brief detail extensive none
Transmit Interval	Time interval, in seconds, used by the transmitting system to send BFD control packets.	brief detail extensive none
Multiplier	Negotiated multiplier by which the time interval is multiplied to determine the detection time for the transmitting system.	detail extensive
Session up time	How long a BFD session has been established.	detail extensive

Table 15: show bfd session Output Fields (*continued*)

Field Name	Field Description	Level of Output
Client	Protocol for which the BFD session is active: ISIS , OSPF , or Static .	detail extensive
TX interval	Time interval, in seconds, used by the host system to transmit BFD control packets.	brief detail extensive none
RX interval	Time interval, in seconds, used by the host system to receive BFD control packets.	brief detail extensive none
Authenticate	Indicates that BFD authentication is configured.	detail extensive
keychain	Name of the security authentication keychain being used by a specific client. BFD authentication information for a client is provided in a single line and includes the keychain , algo , and mode parameters. Multiple clients can be configured on a BFD session.	extensive
algo	BFD authentication algorithm being used for a specific client: keyed-md5 , keyed-sha-1 , meticulous-keyed-md5 , meticulous-keyed-sha-1 , or simple-password . BFD authentication information for a client is provided in a single line and includes the keychain , algo , and mode parameters. Multiple clients can be configured on a BFD session.	extensive
mode	Level of BFD authentication enforcement being used by a specific client: strict or loose . Strict enforcement indicates that authentication is configured at both ends of the session (the default). Loose enforcement indicates that one end of the session might not be authenticated. BFD authentication information for a client is provided in a single line and includes the keychain , algo , and mode parameters. Multiple clients can be configured on a BFD session.	extensive
Local diagnostic	Local diagnostic information about failing BFD sessions.	detail extensive
Remote diagnostic	Remote diagnostic information about failing BFD sessions.	detail extensive
Remote state	Reports whether the remote system's BFD packets have been received and whether the remote system is receiving transmitted control packets.	detail extensive
Version	BFD version: 0 or 1 .	extensive
Replicated	The replicated flag appears when nonstop routing or graceful Routing Engine switchover is configured and the BFD session has been replicated to the backup Routing Engine.	detail extensive
Min async interval	Minimum amount of time, in seconds, between asynchronous control packet transmissions across the BFD session.	extensive
Min slow interval	Minimum amount of time, in seconds, between synchronous control packet transmissions across the BFD session.	extensive

Table 15: show bfd session Output Fields (*continued*)

Field Name	Field Description	Level of Output
Adaptive async TX interval	Transmission interval being used because of adaptation.	extensive
RX interval	Minimum required receive interval.	extensive
Local min TX interval	Minimum amount of time, in seconds, between control packet transmissions on the local system.	extensive
Local min RX interval	Minimum amount of time, in seconds, between control packet detections on the local system.	extensive
Remote min TX interval	Minimum amount of time, in seconds, between control packet transmissions on the remote system.	extensive
Remote min RX interval	Minimum amount of time, in seconds, between control packet detections on the remote system.	extensive
Threshold transmission interval	Threshold for notification if the transmission interval increases.	extensive
Threshold for detection time	Threshold for notification if the detection time increases.	extensive
Local discriminator	Authentication code used by the local system to identify that BFD session.	extensive
Remote discriminator	Authentication code used by the remote system to identify that BFD session.	extensive
Echo mode	Information about the state of echo transmissions on the BFD session.	extensive
Prefix	LDP FEC address associated with the BFD session.	All levels
Egress, Destination	Displays the LDP FEC destination address. This field is displayed only on a router at the egress of an LDP FEC, where the BFD session has an LDP Operation, Administration, and Maintenance (OAM) client.	All levels
Remote is control-plane independent	<p>The BFD session on the remote peer is running on its Packet Forwarding Engine. In this case, when the remote node undergoes a graceful restart, the local peer can help the remote peer with the graceful restart.</p> <p>The following BFD sessions are not distributed to the Packet Forwarding Engine: tunnel-encapsulated sessions, and sessions over integrated routing and bridging (IRB) interfaces.</p>	extensive

Table 15: show bfd session Output Fields (*continued*)

Field Name	Field Description	Level of Output
Authentication	<p>Summary status of BFD authentication:</p> <ul style="list-style-type: none"> status—enabled/active indicates authentication is configured and active. enabled/inactive indicates authentication is configured but not active. This only occurs when the remote end of the session does not support authentication and loose checking is configured. keychain—Name of the security authentication keychain associated with the specified BFD session. algo—BFD authentication algorithm being used: keyed-md5, keyed-sha-1, meticulous-keyed-md5, meticulous-keyed-sha-1, or simple-password. mode—Level of BFD authentication enforcement: strict or loose. Strict enforcement indicates authentication is configured at both ends of the session (the default). Loose enforcement indicates that one end of the session might not be authenticated. <p>This information is only shown if BFD authentication is configured.</p>	extensive
Session ID	The BFD session ID number that represents the protection using MPLS fast reroute (FRR) and loop-free alternate (LFA).	detail extensive
sessions	Total number of active BFD sessions.	All levels
clients	Total number of clients that are hosting active BFD sessions.	All levels
Cumulative transmit rate	Total number of BFD control packets transmitted per second on all active sessions.	All levels
Cumulative receive rate	Total number of BFD control packets received per second on all active sessions.	All levels
Multi-hop, min-recv-TTL	Minimum time to live (TTL) accepted if the session is configured for multihop.	extensive
route table	Route table used if the session is configured for multihop.	extensive
local address	<p>Local address of the source used if the session is configured for multihop.</p> <p>The source IP address for outgoing BFD packets from the egress side of an MPLS BFD session is based on the outgoing interface IP address.</p>	extensive

Sample Output

show bfd session

```
user@host> show bfd session
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.9.1.33	Up	so-7/1/0.0	0.600	0.200	3
10.9.1.29	Up	ge-4/0/0.0	0.600	0.200	3

```
2 sessions, 2 clients
```

```
Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps
```


show bfd session brief

The output for the **show bfd session brief** command is identical to that for the **show bfd session** command. For sample output, see [show bfd session on page 302](#).

show bfd session detail

```
user@host> show bfd session detail
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.9.1.33	Up	so-7/1/0.0	0.600	0.200	3
Client OSPF, TX interval 0.200, RX interval 0.200, multiplier 3					
Session up time 3d 00:34					
Local diagnostic None, remote diagnostic None					
Remote state Up, version 1					
Replicated					
10.9.1.29	Up	ge-4/0/0.0	0.600	0.200	3
Client ISIS L2, TX interval 0.200, RX interval 0.200, multiplier 3					
Session up time 3d 00:29, previous down time 00:00:01					
Local diagnostic NbrSignal, remote diagnostic AdminDown					
Remote state Up, version 1					

2 sessions, 2 clients

Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps

show bfd session detail (with Authentication)

```
user@host> show bfd session detail
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.9.1.33	Up	so-7/1/0.0	0.600	0.200	3
Client OSPF, TX interval 0.200, RX interval 0.200, multiplier 3, Authenticate					
Session up time 3d 00:34					
Local diagnostic None, remote diagnostic None					
Remote state Up, version 1					
Replicated					
10.9.1.29	Up	ge-4/0/0.0	0.600	0.200	3
Client ISIS L2, TX interval 0.200, RX interval 0.200, multiplier 3					
Session up time 3d 00:29, previous down time 00:00:01					
Local diagnostic NbrSignal, remote diagnostic AdminDown					
Remote state Up, version 1					

2 sessions, 2 clients

Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps

show bfd session address extensive

```
user@host> show bfd session 10.255.245.212 extensive
```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.255.245.212	Up		1.200	0.400	3
Client Static, TX interval 0.400, RX interval 0.400, multiplier 3					
Session up time 00:17:03, previous down time 00:00:14					
Local diagnostic CtlExpire, remote diagnostic NbrSignal					
Remote state Up, version 1					
Replicated					
Min async interval 0.400, min slow interval 1.000					
Adaptive async tx interval 0.400, rx interval 0.400					
Local min tx interval 0.400, min rx interval 0.400, multiplier 3					
Remote min tx interval 0.400, min rx interval 0.400, multiplier 3					

Threshold transmission interval 0.000, Threshold for detection time 0.000
 Local discriminator 6, remote discriminator 16
 Echo mode disabled/inactive
 Multi-hop, min-recv-TTL 255, route-table 0, local-address 10.255.245.205

1 sessions, 1 clients
 Cumulative transmit rate 2.5 pps, cumulative receive rate 2.5 pps

show bfd session client rsvp-oam

user@host> show bfd session client rsvp-oam

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
192.168.0.223	Up		540.000	180.000	3

1 Up sessions, 0 Down sessions
 1 sessions, 1 clients
 Cumulative transmit rate 0.0 pps, cumulative receive rate 0.0 pps

show bfd session client vpls-oam summary

user@host> show bfd session client vpls-oam summary

1 Up sessions, 1 Down sessions
 2 sessions, 2 clients
 Cumulative transmit rate 2.0 pps, cumulative receive rate 1.0 pps

show bfd session client vpls-oam instance instance-name

user@host> show bfd session client vpls-oam instance vpls

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
127.0.0.1	Up	ae9.0	3.000	1.000	3

1 Up Sessions, 0 Down Sessions
 1 sessions, 1 clients
 Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps

show bfd session extensive

user@host> show bfd session extensive

10.31.1.2 Up ge-2/1/8.0 0.030 0.010 3
 Client OSPF realm ospf-v2 Area 0.0.0.0, TX interval 0.010, RX interval 0.010
 Session up time 00:10:13
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 Replicated
 Min async interval 0.010, min slow interval 1.000
 Adaptive async TX interval 0.010, RX interval 0.010
 Local min TX interval 0.010, minimum RX interval 0.010, multiplier 3
 Remote min TX interval 0.010, min RX interval 0.010, multiplier 3
 Local discriminator 12, remote discriminator 4
 Echo mode disabled/inactive
 Remote is control-plane independent
 Session ID: 0x201
 Micro-BFD Session

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
10.31.2.2	Up	ge-2/1/4.0	0.030	0.010	3

Client OSPF realm ospf-v2 Area 0.0.0.0, TX interval 0.010, RX interval 0.010

```

Session up time 00:10:14
Local diagnostic None, remote diagnostic NbrSignal
Remote state Up, version 1
Replicated
Min async interval 0.010, min slow interval 1.000
Adaptive async TX interval 0.010, RX interval 0.010
Local min TX interval 0.010, minimum RX interval 0.010, multiplier 3
Remote min TX interval 0.010, min RX interval 0.010, multiplier 3
Local discriminator 13, remote discriminator 5
Echo mode disabled/inactive
Remote is control-plane independent
Session ID: 0x202

```

```

2 sessions, 2 clients
Cumulative transmit rate 200.0 pps, cumulative receive rate 200.0 pps

```

show bfd session extensive (with Authentication)

```

user@host> show bfd session extensive

```

Address	State	Interface	Detect Time	Transmit Interval	Multiplier
192.168.208.26	Up	so-1/0/0.0	2.400	0.800	10

```

Client Static, TX interval 0.600, RX interval 0.600, Authenticate
    keychain bfd, algo keyed-md5, mode loose
Session up time 00:18:07
Local diagnostic None, remote diagnostic NbrSignal
Remote state Up, version 1
Replicated
Min async interval 0.600, min slow interval 1.000
Adaptive async TX interval 0.600, RX interval 0.600
Local min TX interval 0.600, minimum RX interval 0.600, multiplier 10
Remote min TX interval 0.800, min RX interval 0.800, multiplier 3
Local discriminator 2, remote discriminator 3
Echo mode disabled/inactive
Authentication enabled/active, keychain bfd, algo keyed-md5, mode loose

1 sessions, 1 clients
Cumulative transmit rate 1.2 pps, cumulative receive rate 1.2 pps

```

show bfd session summary

```

user@host> show bfd session summary
2 sessions, 2 clients
Cumulative transmit rate 10.0 pps, cumulative receive rate 10.0 pps

```


CHAPTER 8

BGP Operational Mode Commands

- clear bgp damping
- clear bgp neighbor
- clear bgp table
- show bgp bmp
- show bgp group
- show bgp group traffic-statistics
- show bgp neighbor
- show bgp replication
- show bgp summary
- show policy damping

clear bgp damping

Syntax	clear bgp damping <logical-system (all <i>logical-system-name</i>)> < <i>prefix</i> >
Syntax (EX Series Switch and QFX Series)	clear bgp damping < <i>prefix</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Clear BGP route flap damping information.
Options	none —Clear all BGP route flap damping information. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. <i>prefix</i> —(Optional) Clear route flap damping information for only the specified destination prefix.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show policy damping on page 347• show route damping on page 802
List of Sample Output	clear bgp damping on page 308
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear bgp damping

```
user@host> clear bgp damping
```

clear bgp neighbor

Syntax	<pre>clear bgp neighbor <as <i>as-number</i>> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <malformed-route> <neighbor> <soft soft-inbound> <soft-minimum-igp></pre>
Syntax (EX Series Switch and QFX Series)	<pre>clear bgp neighbor <as <i>as-number</i>> <instance <i>instance-name</i>> <malformed-route> <neighbor> <soft soft-inbound> <soft-minimum-igp></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>malformed-route option introduced in Junos OS Release 13.2.</p>
Description	<p>Perform one of the following tasks:</p> <ul style="list-style-type: none"> • Change the state of one or more BGP neighbors to IDLE. For neighbors in the ESTABLISHED state, this command drops the TCP connection to the neighbors and then reestablishes the connection. • (soft keyword only) Reapply export policies or import policies, respectively, to one or more BGP neighbors without changing their state. • (soft-inbound keyword only) Reapply export policies or import policies, respectively, and send refresh updates to one or more BGP neighbors without changing their state.
Options	<p>none—Change the state of all BGP neighbors to IDLE.</p> <p>as <i>as-number</i>—(Optional) Apply this command only to neighbors in the specified autonomous system (AS).</p> <p>instance <i>instance-name</i>—(Optional) Apply this command only to neighbors for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>malformed-route—(Optional) Remove malformed routes. If a specific neighbor is provided, Junos OS removes malformed routes for that particular neighbor. Otherwise, Junos OS removes malformed routes for all BGP neighbors. To find routes that have malformed attributes, run the show route hidden command, and look for routes marked with MalformedAttr in the AS path field.</p>

neighbor—(Optional) IP address of a BGP peer. Apply this command only to the specified neighbor.

soft—(Optional) Reapply any export policies and send refresh updates to neighbors without clearing the state.

soft-inbound—(Optional) Reapply any import policies and send refresh updates to neighbors without clearing the state.

soft-minimum-igp—(Optional) Provides soft refresh of the outbound state when the interior gateway protocol (IGP) metric is reset.

Required Privilege Level

clear

Related Documentation

- [show bgp neighbor on page 324](#)

List of Sample Output

[clear bgp neighbor on page 310](#)

Output Fields

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

Sample Output

`clear bgp neighbor`

```
user@host> clear bgp neighbor
```


clear bgp table

Syntax	<code>clear bgp table <i>table-name</i></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switch and QFX Series)	<code>clear bgp table <i>table-name</i></code>
Release Information	Command introduced in Junos OS Release 9.0. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Request that BGP refresh routes in a specified routing table.
Options	<code>logical-system (all <i>logical-system-name</i>)</code> —(Optional) Perform this operation on all logical systems or on a particular logical system. <code>table-name</code> —Request that BGP refresh routes in the specified table.
Additional Information	In some cases, a prefix limit is associated with a routing table for a VPN instance. When this limit is exceeded (for example, because of a network misconfiguration), some routes might not be inserted in the table. Such routes need to be added to the table after the network issue is resolved. Use the clear bgp table command to request that BGP refresh routes in a VPN instance table.
Required Privilege Level	clear
List of Sample Output	clear bgp table private.inet.0 on page 311 clear bgp table inet.6 logical-system all on page 311 clear bgp table private.inet.6 logical-system ls1 on page 311 clear bgp table logical-system all inet.0 on page 311 clear bgp table logical-system ls2 private.inet.0 on page 312
Output Fields	This command produces no output.

Sample Output

`clear bgp table private.inet.0`

```
user@host> clear bgp table private.inet.0
```

`clear bgp table inet.6 logical-system all`

```
user@host> clear bgp table inet.6 logical-system all
```

`clear bgp table private.inet.6 logical-system ls1`

```
user@host> clear bgp table private.inet.6 logical-system ls1
```

`clear bgp table logical-system all inet.0`

```
user@host> clear bgp table logical-system all inet.0
```

clear bgp table logical-system ls2 private.inet.0

```
user@host> clear bgp table logical-system ls2 private.inet.0
```

show bgp bmp

Syntax	show bgp bmp
Release Information	Command introduced in Junos OS Release 9.5. Command introduced in Junos OS Release 9.5 for EX Series switches. Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Description	Display information about the BGP Monitoring Protocol (BMP).
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show bgp bmp on page 313
Output Fields	Table 16 on page 313 lists the output fields for the show bgp bmp command. Output fields are listed in the approximate order in which they appear.

Table 16: show bgp bmp Output Fields

Field Name	Field Description
BMP station address/port	IP address and port number of the monitoring station to which BGP Monitoring Protocol (BMP) statistics are sent.
BMP session state	Status of the BMP session: UP or DOWN .
Memory consumed by BMP	Memory used by the active BMP session.
Statistics timeout	Amount of time, in seconds, between transmissions of BMP data to the monitoring station.
Memory limit	Threshold, in bytes, at which the routing device stops collecting BMP data.
Memory-connect retry timeout	Amount of time, in seconds, after which the routing device attempts to resume a BMP session that was ended after the configured memory threshold was exceeded.

Sample Output

show bgp bmp

```

user@host> show bgp bmp
  BMP station address/port: 172.24.24.157+5454
  BMP session state: DOWN
  Memory consumed by BMP: 0
  Statistics timeout: 15
  Memory limit: 10485760
  Memory connect retry timeout: 600

```


show bgp group

Syntax	<pre>show bgp group <brief detail summary> <group-name> <exact-instance instance-name> <instance instance-name> <logical-system (all logical-system-name)> <rtf></pre>
Syntax (EX Series Switch and QFX Series)	<pre>show bgp group <brief detail summary> <group-name> <exact-instance instance-name> <instance instance-name></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>exact-instance option introduced in Junos OS Release 11.4.</p>
Description	Display information about the configured BGP groups.
Options	<p>none—Display group information about all BGP groups.</p> <p>brief detail summary—(Optional) Display the specified level of output.</p> <p>group-name—(Optional) Display group information for the specified group.</p> <p>exact-instance instance-name—(Optional) Display information for the specified instance only.</p> <p>instance instance-name—(Optional) Display information about BGP groups for all routing instances whose name begins with this string (for example, cust1, cust11, and cust111 are all displayed when you run the show bgp group instance cust1 command). The instance name can be master for the main instance, or any valid configured instance name or its prefix.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>rtf—(Optional) Display BGP group route targeting information.</p>
Required Privilege Level	view
List of Sample Output	<p>show bgp group on page 319</p> <p>show bgp group brief on page 319</p> <p>show bgp group detail on page 320</p> <p>show bgp group rtf detail on page 321</p> <p>show bgp group summary on page 321</p>

Output Fields Table 17 on page 316 describes the output fields for the **show bgp group** command. Output fields are listed in the approximate order in which they appear.

Table 17: show bgp group Output Fields

Field Name	Field Description	Level of Output
Group Type or Group	Type of BGP group: Internal or External .	All levels
group-index	Index number for the BGP peer group. The index number differentiates between groups when a single BGP group is split because of different configuration options at the group and peer levels.	rtf detail
AS	AS number of the peer. For internal BGP (IBGP), this number is the same as Local AS .	brief detail none
Local AS	AS number of the local routing device.	brief detail none
Name	Name of a specific BGP group.	brief detail none
Index	Unique index number of a BGP group.	brief detail none
Flags	Flags associated with the BGP group. This field is used by Juniper Networks customer support.	brief detail none
Remove-private options	Options associated with the remove-private statement.	brief detail none
Holdtime	Maximum number of seconds allowed to elapse between successive keepalive or update messages that BGP receives from a peer in the BGP group, after which the connection to the peer is closed and routing devices through that peer become unavailable.	brief detail none
Export	Export policies configured for the BGP group with the export statement.	brief detail none
MED tracks IGP metric update delay	Time, in seconds, that updates to multiple exit discriminator (MED) are delayed. Also displays the time remaining before the interval is set to expire	All levels
Traffic Statistics Interval	Time between sample periods for labeled-unicast traffic statistics, in seconds.	brief detail none
Total peers	Total number of peers in the group.	brief detail none
Established	Number of peers in the group that are in the established state.	All levels

Table 17: show bgp group Output Fields (*continued*)

Field Name	Field Description	Level of Output
Active/Received/Accepted/Damped	<p>Multipurpose field that displays information about BGP peer sessions. The field's contents depend upon whether a session is established and whether it was established in the main routing device or in a routing instance.</p> <ul style="list-style-type: none"> If a peer is not established, the field shows the state of the peer session: Active, Connect, or Idle. If a BGP session is established in the main routing device, the field shows the number of active, received, accepted, and damped routes that are received from a neighbor and appear in the inet.0 (main) and inet.2 (multicast) routing tables. For example, 8/10/10/2 and 2/4/4/0 indicate the following: <ul style="list-style-type: none"> 8 active routes, 10 received routes, 10 accepted routes, and 2 damped routes from a BGP peer appear in the inet.0 routing table. 2 active routes, 4 received routes, 4 accepted routes, and no damped routes from a BGP peer appear in the inet.2 routing table. 	summary
ip-addresses	List of peers who are members of the group. The address is followed by the peer's port number.	All levels
Route Queue Timer	Number of seconds until queued routes are sent. If this time has already elapsed, this field displays the number of seconds by which the updates are delayed.	detail
Route Queue	Number of prefixes that are queued up for sending to the peers in the group.	detail
inet.number	<p>Number of active, received, accepted, and damped routes in the routing table. For example, inet.0: 7/10/9/0 indicates the following:</p> <ul style="list-style-type: none"> 7 active routes, 10 received routes, 9 accepted routes, and no damped routes from a BGP peer appear in the inet.0 routing table. 	none

Table 17: show bgp group Output Fields (*continued*)

Field Name	Field Description	Level of Output
Table inet.number	<p>Information about the routing table.</p> <ul style="list-style-type: none"> • Received prefixes—Total number of prefixes from the peer, both active and inactive, that are in the routing table. • Active prefixes—Number of prefixes received from the peer that are active in the routing table. • Suppressed due to damping—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols. • Advertised prefixes—Number of prefixes advertised to a peer. • Received external prefixes—Total number of prefixes from the external BGP (EBGP) peers, both active and inactive, that are in the routing table. • Active external prefixes—Number of prefixes received from the EBGP peers that are active in the routing table. • Externals suppressed—Number of routes received from EBGP peers currently inactive because of damping or other reasons. • Received internal prefixes—Total number of prefixes from the IBGP peers, both active and inactive, that are in the routing table. • Active internal prefixes—Number of prefixes received from the IBGP peers that are active in the routing table. • Internals suppressed—Number of routes received from IBGP peers currently inactive because of damping or other reasons. • RIB State—Status of the graceful restart process for this routing table: BGP restart is complete, BGP restart in progress, VPN restart in progress, or VPN restart is complete. 	detail
Groups	Total number of groups.	All levels
Peers	Total number of peers.	All levels
External	Total number of external peers.	All levels
Internal	Total number of internal peers.	All levels
Down peers	Total number of unavailable peers.	All levels
Flaps	Total number of flaps that occurred.	All levels
Table	Name of a routing table.	brief , none
Tot Paths	Total number of routes.	brief , none
Act Paths	Number of active routes.	brief , none
Suppressed	Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.	brief , none

Table 17: show bgp group Output Fields (*continued*)

Field Name	Field Description	Level of Output
History	Number of withdrawn routes stored locally to keep track of damping history.	brief, none
Damp State	Number of active routes with a figure of merit greater than zero, but lower than the threshold at which suppression occurs.	brief, none
Pending	Routes being processed by the BGP import policy.	brief, none
Group	Group the peer belongs to in the BGP configuration.	detail
Receive mask	Mask of the received target included in the advertised route.	detail
Entries	Number of route entries received.	detail
Target	Route target that is to be passed by route-target filtering. If a route advertised from the provider edge (PE) routing device matches an entry in the route-target filter, the route is passed to the peer.	detail
Mask	Mask which specifies that the peer receive routes with the given route target.	detail

Sample Output

show bgp group

```

user@host> show bgp group
Groups: 2  Peers: 2   External: 0   Internal: 2   Down peers: 1   Flaps: 0
Table      Tot Paths  Act Paths  Suppressed   History Damp State   Pending

inet.0
          0         0         0           0         0         0

bgp.13vpn.0
          0         0         0           0         0         0

bgp.rtarget.0
          2         0         0           0         0         0

```

show bgp group brief

```

user@host> show bgp group brief
Groups: 2  Peers: 2   External: 0   Internal: 2   Down peers: 1   Flaps: 0
Table      Tot Paths  Act Paths  Suppressed   History Damp State   Pending

inet.0
          0         0         0           0         0         0

bgp.13vpn.0
          0         0         0           0         0         0

bgp.rtarget.0
          2         0         0           0         0         0

```

show bgp group detail

```

user@host> show bgp group detail
Group Type: Internal   AS: 1                      Local AS: 1
Name: ibgp             Index: 0                    Flags: <Export Eval>
Holdtime: 0
Total peers: 3         Established: 0
22.0.0.2
22.0.0.8
22.0.0.5

Groups: 1 Peers: 3   External: 0   Internal: 3   Down peers: 3   Flaps: 3
Table bgp.l3vpn.0
  Received prefixes:      0
  Accepted prefixes:      0
  Active prefixes:        0
  Suppressed due to damping: 0
  Received external prefixes: 0
  Active external prefixes: 0
  Externals suppressed:   0
  Received internal prefixes: 0
  Active internal prefixes: 0
  Internals suppressed:   0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
Table bgp.mdt.0
  Received prefixes:      0
  Accepted prefixes:      0
  Active prefixes:        0
  Suppressed due to damping: 0
  Received external prefixes: 0
  Active external prefixes: 0
  Externals suppressed:   0
  Received internal prefixes: 0
  Active internal prefixes: 0
  Internals suppressed:   0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
Table VPN-A.inet.0
  Received prefixes:      0
  Accepted prefixes:      0
  Active prefixes:        0
  Suppressed due to damping: 0
  Received external prefixes: 0
  Active external prefixes: 0
  Externals suppressed:   0
  Received internal prefixes: 0
  Active internal prefixes: 0
  Internals suppressed:   0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
Table VPN-A.mdt.0
  Received prefixes:      0
  Accepted prefixes:      0
  Active prefixes:        0
  Suppressed due to damping: 0
  Received external prefixes: 0
  Active external prefixes: 0
  Externals suppressed:   0
  Received internal prefixes: 0
  Active internal prefixes: 0

```

```

Internals suppressed:      0
RIB State: BGP restart is complete
RIB State: VPN restart is complete

```

show bgp group rtf detail

```

user@host> show bgp group rtf detail
Group: internal (group-index: 0)
  Receive mask: 00000002
  Table: bgp.rtarget.0
    Target
    100:100/64
    200:201/64
    Mask
    00000002
    (Group)
  Entries: 2
Group: internal (group-index: 1)
  Table: bgp.rtarget.0
    Target
    200:201/64
    Mask
    (Group)
  Entries: 1

```

show bgp group summary

```

user@host> show bgp group summary
Group      Type      Peers      Established      Active/Received/Accepted/Damped
ibgp       Internal  3          0
Groups: 1  Peers: 3      External: 0      Internal: 3      Down peers: 3      Flaps: 3
bgp.l3vpn.0 : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
bgp.mdt.0   : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
VPN-A.inet.0 : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
VPN-A.mdt.0 : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0

```

show bgp group traffic-statistics

Syntax	show bgp group traffic-statistics <brief detail> <group-name> <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display the traffic statistics for configured Border Gateway Protocol (BGP) groups.
Options	<p>none—Display traffic statistics for all BGP groups.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>group-name—(Optional) Display BGP traffic statistics for only the specified group.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show bgp group traffic-statistics (Per-Group-Label Not Configured) on page 323 show bgp group traffic-statistics (Per-Group-Label Configured) on page 323
Output Fields	Table 18 on page 322 describes the output fields for the show bgp group traffic-statistics command. Output fields are listed in the approximate order in which they appear.

Table 18: show bgp group traffic-statistics Output Fields

Field Name	Field Description
Group name	Name of a specific BGP group.
Group Index	Index number for the BGP group.
NLRI	Network layer reachability information (NLRI) indicating the source of the traffic statistics for the BGP group.
FEC	Forwarding equivalence classes (FECs) associated with the BGP group.
Packets	Number of packets sent through each FEC.
Bytes	Number of bytes transmitted through each FEC.
EgressAS	Autonomous system (AS) number of the egress router.
AdvLabel	Label associated with each FEC.

Sample Output

show bgp group traffic-statistics (Per-Group-Label Not Configured)

```

user@host> show bgp group traffic-statistics
Group Name: ext1      Group Index: 0      NLRI: inet-labeled-unicast
FEC                   Packets            Bytes      EgressAS    AdvLabel
10.255.245.55         0                  0          I           100224
10.255.245.57         0                  0          I           100240
100.101.0.0           550                48400      25          100256
100.102.0.0           550                48400      25          100256
100.103.0.0           550                48400      25          100272
100.104.0.0           550                48400      25          100272
192.168.25.0          0                  0          I           100288

Group Name: ext2      Group Index: 1      NLRI: inet-labeled-unicast
FEC                   Packets            Bytes      EgressAS    AdvLabel
10.255.245.55         0                  0          I           100224
10.255.245.57         0                  0          I           100240
100.101.0.0           550                48400      25          100256
100.102.0.0           550                48400      25          100256
100.103.0.0           550                48400      25          100272
100.104.0.0           550                48400      25          100272
192.168.25.0          0                  0          I           100288

```

show bgp group traffic-statistics (Per-Group-Label Configured)

```

user@host> show bgp group traffic-statistics
Group Name: ext1      Group Index: 0      NLRI: inet-labeled-unicast
FEC                   Packets            Bytes      EgressAS    AdvLabel
10.255.245.55         0                  0          I           100384
10.255.245.57         0                  0          I           100400
100.101.0.0           101                8888       25          100416
100.102.0.0           101                8888       25          100416
100.103.0.0           0                  0          25          100432
100.104.0.0           0                  0          25          100432
192.168.25.0          0                  0          I           100448

Group Name: ext2      Group Index: 1      NLRI: inet-labeled-unicast
FEC                   Packets            Bytes      EgressAS    AdvLabel
10.255.245.55         0                  0          I           100304
10.255.245.57         0                  0          I           100320
100.101.0.0           0                  0          25          100336
100.102.0.0           0                  0          25          100336
100.103.0.0           101                8888       25          100352
100.104.0.0           101                8888       25          100352
192.168.25.0          0                  0          I           100368

```

show bgp neighbor

Syntax	<pre>show bgp neighbor <exact-instance <i>instance-name</i>> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <neighbor-address> <orf (detail <i>neighbor-address</i>)</pre>
Syntax (EX Series Switch and QFX Series)	<pre>show bgp neighbor <instance <i>instance-name</i>> <exact-instance <i>instance-name</i>> <neighbor-address> <orf (<i>neighbor-address</i> detail)</pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>orf option introduced in Junos OS Release 9.2.</p> <p>exact-instance option introduced in Junos OS Release 11.4.</p>
Description	Display information about BGP peers.
Options	<p>none—Display information about all BGP peers.</p> <p>exact-instance <i>instance-name</i>—(Optional) Display information for the specified instance only.</p> <p>instance <i>instance-name</i>—(Optional) Display information about BGP peers for all routing instances whose name begins with this string (for example, cust1, cust11, and cust111 are all displayed when you run the show bgp neighbor instance cust1 command).</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>neighbor-address—(Optional) Display information for only the BGP peer at the specified IP address.</p> <p>orf (detail <i>neighbor-address</i>)—(Optional) Display outbound route-filtering information for all BGP peers or only for the BGP peer at the specified IP address. The default is to display brief output. Use the detail option to display detailed output.</p>
Additional Information	For information about the local-address , nlri , hold-time , and preference statements, see the <i>Junos OS Routing Protocols Library for Routing Devices</i> .
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• clear bgp neighbor on page 309

List of Sample Output [show bgp neighbor on page 331](#)
[show bgp neighbor \(CLNS\) on page 332](#)
[show bgp neighbor \(Layer 2 VPN\) on page 333](#)
[show bgp neighbor \(Layer 3 VPN\) on page 335](#)
[show bgp neighbor neighbor-address on page 335](#)
[show bgp neighbor neighbor-address on page 336](#)
[show bgp neighbor orf neighbor-address detail on page 337](#)

Output Fields [Table 19 on page 325](#) describes the output fields for the **show bgp neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 19: show bgp neighbor Output Fields

Field Name	Field Description
Peer	Address of the BGP neighbor. The address is followed by the neighbor port number.
AS	AS number of the peer.
Local	Address of the local routing device. The address is followed by the peer port number.
Type	Type of peer: Internal or External .
State	<p>Current state of the BGP session:</p> <ul style="list-style-type: none"> • Active—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message. • Connect—BGP is waiting for the transport protocol connection to be completed. • Established—The BGP session has been established, and the peers are exchanging update messages. • Idle—This is the first stage of a connection. BGP is waiting for a Start event. • OpenConfirm—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message. • OpenSent—BGP has sent an open message and is waiting to receive an open message from the peer.
Flags	<p>Internal BGP flags:</p> <ul style="list-style-type: none"> • Aggregate Label—BGP has aggregated a set of incoming labels (labels received from the peer) into a single forwarding label. • CleanUp—The peer session is being shut down. • Delete—This peer has been deleted. • Idled—This peer has been permanently idled. • ImportEval—At the last commit operation, this peer was identified as needing to reevaluate all received routes. • Initializing—The peer session is initializing. • SendRtn—Messages are being sent to the peer. • Sync—This peer is synchronized with the rest of the peer group. • TryConnect—Another attempt is being made to connect to the peer. • Unconfigured—This peer is not configured. • WriteFailed—An attempt to write to this peer failed.

Table 19: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Last state	<p>Previous state of the BGP session:</p> <ul style="list-style-type: none"> • Active—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message. • Connect—BGP is waiting for the transport protocol connection to be completed. • Established—The BGP session has been established, and the peers are exchanging update messages. • Idle—This is the first stage of a connection. BGP is waiting for a Start event. • OpenConfirm—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message. • OpenSent—BGP has sent an open message and is waiting to receive an open message from the peer.
Last event	<p>Last activity that occurred in the BGP session:</p> <ul style="list-style-type: none"> • Closed—The BGP session closed. • ConnectRetry—The transport protocol connection failed, and BGP is trying again to connect. • HoldTime—The session ended because the hold timer expired. • KeepAlive—The local routing device sent a BGP keepalive message to the peer. • Open—The local routing device sent a BGP open message to the peer. • OpenFail—The local routing device did not receive an acknowledgment of a BGP open message from the peer. • RecvKeepAlive—The local routing device received a BGP keepalive message from the peer. • RecvNotify—The local routing device received a BGP notification message from the peer. • RecvOpen—The local routing device received a BGP open message from the peer. • RecvUpdate—The local routing device received a BGP update message from the peer. • Start—The peering session started. • Stop—The peering session stopped. • TransportError—A TCP error occurred.
Last error	<p>Last error that occurred in the BGP session:</p> <ul style="list-style-type: none"> • Cease—An error occurred, such as a version mismatch, that caused the session to close. • Finite State Machine Error—In setting up the session, BGP received a message that it did not understand. • Hold Time Expired—The session's hold time expired. • Message Header Error—The header of a BGP message was malformed. • Open Message Error—A BGP open message contained an error. • None—No errors occurred in the BGP session. • Update Message Error—A BGP update message contained an error.
Export	Name of the export policy that is configured on the peer.
Import	Name of the import policy that is configured on the peer.

Table 19: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Options	Configured BGP options: <ul style="list-style-type: none"> • AddressFamily—Configured address family: inet or inet-vpn. • AuthKeyChain—Authentication key change is enabled. • DropPathAttributes—Certain path attributes are configured to be dropped from neighbor updates during inbound processing. • GracefulRestart—Graceful restart is configured. • HoldTime—Hold time configured with the hold-time statement. The hold time is three times the interval at which keepalive messages are sent. • IgnorePathAttributes—Certain path attributes are configured to be ignored in neighbor updates during inbound processing. • Local Address—Address configured with the local-address statement. • Multihop—Allow BGP connections to external peers that are not on a directly connected network. • NLRI—Configured MBGP state for the BGP group: multicast, unicast, or both if you have configured nlri any. • Peer AS—Configured peer autonomous system (AS). • Preference—Preference value configured with the preference statement. • Refresh—Configured to refresh automatically when the policy changes. • Rib-group—Configured routing table group.
Path-attributes dropped	Path attribute codes that are dropped from neighbor updates.
Path-attributes ignored	Path attribute codes that are ignored during neighbor updates.
Authentication key change	(appears only if the authentication-keychain statement has been configured) Name of the authentication keychain enabled.
Authentication algorithm	(appears only if the authentication-algorithm statement has been configured) Type of authentication algorithm enabled: hmac or md5 .
Address families configured	Names of configured address families for the VPN.
Local Address	Address of the local routing device.
Remove-private options	Options associated with the remove-private statement.
Holdtime	Hold time configured with the hold-time statement. The hold time is three times the interval at which keepalive messages are sent.
Flags for NLRI inet-label-unicast	Flags related to labeled-unicast: <ul style="list-style-type: none"> • TrafficStatistics—Collection of statistics for labeled-unicast traffic is enabled.

Table 19: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Traffic statistics	Information about labeled-unicast traffic statistics: <ul style="list-style-type: none"> • Options—Options configured for collecting statistics about labeled-unicast traffic. • File—Name and location of statistics log files. • size—Size of all the log files, in bytes. • files—Number of log files.
Traffic Statistics Interval	Time between sample periods for labeled-unicast traffic statistics, in seconds.
Preference	Preference value configured with the preference statement.
Outbound Timer	Time for which the route is available in Junos OS routing table before it is exported to BGP. This field is displayed in the output only if the out-delay parameter is configured to a non-zero value.
Number of flaps	Number of times the BGP session has gone down and then come back up.
Peer ID	Router identifier of the peer.
Group index	Index number for the BGP peer group. The index number differentiates between groups when a single BGP group is split because of different configuration options at the group and peer levels.
Peer index	Index that is unique within the BGP group to which the peer belongs.
Local ID	Router identifier of the local routing device.
Local Interface	Name of the interface on the local routing device.
Active holdtime	Hold time that the local routing device negotiated with the peer.
Keepalive Interval	Keepalive interval, in seconds.
BFD	Status of BFD failure detection.
Local Address	Name of directly connected interface over which direct EBGP peering is established.
NLRI for restart configured on peer	Names of address families configured for restart.
NLRI advertised by peer	Address families supported by the peer: unicast or multicast .
NLRI for this session	Address families being used for this session.
Peer supports Refresh capability	Remote peer's ability to send and request full route table readvertisement (route refresh capability). For more information, see RFC 2918, <i>Route Refresh Capability for BGP-4</i> .
Restart time configured on peer	Configured time allowed for restart on the neighbor.

Table 19: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Stale routes from peer are kept for	When graceful restart is negotiated, the maximum time allowed to hold routes from neighbors after the BGP session has gone down.
Peer does not support Restarter functionality	Graceful restart restarter-mode is disabled on the peer.
Peer does not support Receiver functionality	Graceful restart helper-mode is disabled on the peer.
Restart time requested by this peer	Restart time requested by this neighbor during capability negotiation.
Restart flag received from the peer	When this field appears, the BGP speaker has restarted (Restarting), and this peer should not wait for the end-of-rib marker from the speaker before advertising routing information to the speaker.
NLRI that peer supports restart for	Neighbor supports graceful restart for this address family.
NLRI peer can save forwarding state	Neighbor supporting this address family saves all forwarding states.
NLRI that peer saved forwarding for	Neighbor saves all forwarding states for this address family.
NLRI that restart is negotiated for	Router supports graceful restart for this address family.
NLRI of received end-of-rib markers	Address families for which end-of-routing-table markers are received from the neighbor.
NLRI of all end-of-rib markers sent	Address families for which end-of-routing-table markers are sent to the neighbor.
Peer supports 4 byte AS extension (peer-as 1)	Peer understands 4-byte AS numbers in BGP messages. The peer is running Junos OS Release 9.1 or later.
NLRIs for which peer can receive multiple paths	Appears in the command output of the local router if the downstream peer is configured to receive multiple BGP routes to a single destination, instead of only receiving the active route. Possible value is inet-unicast .
NLRIs for which peer can send multiple paths: inet-unicast	Appears in the command output of the local router if the upstream peer is configured to send multiple BGP routes to a single destination, instead of only sending the active route. Possible value is inet-unicast .

Table 19: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Table inet. <i>number</i>	<p>Information about the routing table:</p> <ul style="list-style-type: none"> • RIB State—BGP is in the graceful restart process for this routing table: restart is complete or restart in progress. • Bit—Number that represents the entry in the routing table for this peer. • Send state—State of the BGP group: in sync, not in sync, or not advertising. • Active prefixes—Number of prefixes received from the peer that are active in the routing table. • Received prefixes—Total number of prefixes from the peer, both active and inactive, that are in the routing table. • Accepted prefixes—Total number of prefixes from the peer that have been accepted by a routing policy. • Suppressed due to damping—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.
Last traffic (seconds)	Last time any traffic was received from the peer or sent to the peer, and the last time the local routing device checked.
Input messages	Messages that BGP has received from the receive socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB.
Output messages	Messages that BGP has written to the transmit socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB.
Input dropped path attributes	<p>Information about dropped path attributes:</p> <ul style="list-style-type: none"> • Code—Path attribute code. • Count—Path attribute count.
Input ignored path attributes	<p>Information about ignored path attributes:</p> <ul style="list-style-type: none"> • Code—Path attribute code. • Count—Path attribute count.
Output queue	<p>Number of BGP packets that are queued to be transmitted to a particular neighbor for a particular routing table. Output queue 0 is for unicast NLRIs, and queue 1 is for multicast NLRIs.</p> <p>It also specifies the routing table name and the NLRI they represent in the format (<i>routing table name</i>, <i>NLRI</i>).</p> <p>NOTE: The output queues of routing tables that are not advertised, will only show up at extensive output level.</p>
Trace options	Configured tracing of BGP protocol packets and operations.
Trace file	Name of the file to receive the output of the tracing operation.

Table 19: show bgp neighbor Output Fields (*continued*)

Field Name	Field Description
Filter Updates rcv	(orf option only) Number of outbound-route filters received for each configured address family. NOTE: The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list.
Immediate	(orf option only) Number of route updates received with the immediate flag set. The immediate flag indicates that the BGP peer should readvertise the updated routes. NOTE: The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list.
Filter	(orf option only) Type of prefix filter received: prefix-based or extended-community .
Received filter entries	(orf option only) List of received filters displayed.
seq	(orf option only) Numerical order assigned to this prefix entry among all the received outbound route filter prefix entries.
prefix	(orf option only) Address for the prefix entry that matches the filter.
minlength	(orf option only) Minimum prefix length, in bits, required to match this prefix.
maxlength	(orf option only) Maximum prefix length, in bits, required to match this prefix.
match	(orf option only) For this prefix match, whether to permit or deny route updates.

Sample Output

show bgp neighbor

```

user@host > show bgp neighbor
Peer: 10.255.7.250+179 AS 10   Local: 10.255.7.248+63740 AS 10
  Type: Internal   State: Established   Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ redist_static ]
  Options: <Preference LocalAddress PeerAS Refresh>
  Local Address: 10.255.7.248 Holdtime: 90 Preference: 170 Outbound Timer: 50
  Number of flaps: 0
  Peer ID: 10.255.7.250   Local ID: 10.255.7.248   Active Holdtime: 90
  Keepalive Interval: 30   Group index: 0   Peer index: 0
  BFD: disabled, down
  NLRI for restart configured on peer: inet-unicast
  NLRI advertised by peer: inet-unicast
  NLRI for this session: inet-unicast
  Peer supports Refresh capability (2)
  Stale routes from peer are kept for: 300
  Peer does not support Restarter functionality
  NLRI that restart is negotiated for: inet-unicast
  NLRI of received end-of-rib markers: inet-unicast
  NLRI of all end-of-rib markers sent: inet-unicast
  Peer supports 4 byte AS extension (peer-as 10)

```

```

Peer does not support Addpath
Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Accepted prefixes:        1
  Suppressed due to damping: 0
  Advertised prefixes:      1
Last traffic (seconds): Received 9    Sent 5    Checked 5
Input messages:  Total 36    Updates 2    Refreshes 0    Octets 718
Output messages: Total 37    Updates 1    Refreshes 0    Octets 796
Output Queue[0]: 0 (inet.0, inet-unicast)

Peer: 10.255.162.214+52193 AS 100 Local: 10.255.167.205+179 AS 100
  Type: Internal    State: Established (route reflector client)Flags: <Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress Cluster AddressFamily Rib-group Refresh>
  Address families configured: inet-unicast inet-vpn-unicast route-target
  Local Address: 10.255.167.205 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.255.162.214    Local ID: 10.255.167.205    Active Holdtime: 90
  Keepalive Interval: 30    Group index: 0    Peer index: 1

```

show bgp neighbor (CLNS)

```

user@host> show bgp neighbor
Peer: 10.245.245.1+179 AS 200 Local: 10.245.245.3+3770 AS 100
  Type: External    State: Established    Flags: <ImportEval Sync>
  Last State: OpenConfirm    Last Event: RecvKeepAlive
  Last Error: None
  Options: <Multihop Preference LocalAddress HoldTime AddressFamily PeerAS
  Rib-group Refresh>
  Address families configured: iso-vpn-unicast
  Local Address: 10.245.245.3 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.245.245.1    Local ID: 10.245.245.3    Active Holdtime: 90
  Keepalive Interval: 30    Peer index: 0
  NLRI advertised by peer: iso-vpn-unicast
  NLRI for this session: iso-vpn-unicast
  Peer supports Refresh capability (2)
Table bgp.isovpn.0 Bit: 10000
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: in sync
  Active prefixes:          3
  Received prefixes:        3
  Suppressed due to damping: 0
  Advertised prefixes:      3
Table aaaa.iso.0
  RIB State: BGP restart is complete
  RIB State: VPN restart is complete
  Send state: not advertising
  Active prefixes:          3
  Received prefixes:        3
  Suppressed due to damping: 0
Last traffic (seconds): Received 6    Sent 5    Checked 5
Input messages:  Total 1736    Updates 4    Refreshes 0    Octets 33385
Output messages: Total 1738    Updates 3    Refreshes 0    Octets 33305
Output Queue[0]: 0 (bgp.isovpn.0, iso-vpn-unicast)
Output Queue[1]: 0 (aaaa.iso.0, iso-vpn-unicast)

```

show bgp neighbor (Layer 2 VPN)

```

user@host> show bgp neighbor
Peer: 10.69.103.2      AS 65100 Local: 10.69.103.1      AS 65103
  Type: External      State: Active      Flags: <ImportEval>
  Last State: Idle      Last Event: Start
  Last Error: None
  Export: [ BGP-INET-import ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
  Address families configured: inet-unicast
  Local Address: 10.69.103.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
Peer: 10.69.104.2      AS 65100 Local: 10.69.104.1      AS 65104
  Type: External      State: Active      Flags: <ImportEval>
  Last State: Idle      Last Event: Start
  Last Error: None
  Export: [ BGP-L-import ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
Refresh>
  Address families configured: inet-labeled-unicast
  Local Address: 10.69.104.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
Peer: 10.255.14.182+179 AS 69      Local: 10.255.14.176+2131 AS 69
  Type: Internal      State: Established      Flags: <ImportEval>
  Last State: OpenConfirm      Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
Rib-group Refresh>
  Address families configured: inet-vpn-unicast l2vpn
  Local Address: 10.255.14.176 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.255.14.182      Local ID: 10.255.14.176      Active Holdtime: 90
  Keepalive Interval: 30
  NLRI for restart configured on peer: inet-vpn-unicast l2vpn
  NLRI advertised by peer: inet-vpn-unicast l2vpn
  NLRI for this session: inet-vpn-unicast l2vpn
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 120
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-vpn-unicast l2vpn
  NLRI peer can save forwarding state: inet-vpn-unicast l2vpn
  NLRI that peer saved forwarding for: inet-vpn-unicast l2vpn
  NLRI that restart is negotiated for: inet-vpn-unicast l2vpn
  NLRI of received end-of-rib markers: inet-vpn-unicast l2vpn
Table bgp.l3vpn.0 Bit: 10000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          10
  Received prefixes:        10
  Suppressed due to damping: 0
Table bgp.l2vpn.0 Bit: 20000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:          1
  Received prefixes:        1
  Suppressed due to damping: 0
Table BGP-INET.inet.0 Bit: 30000

```

```

RIB State: BGP restart in progress
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:          2
Received prefixes:        2
Suppressed due to damping: 0
Table BGP-L.inet.0 Bit: 40000
RIB State: BGP restart in progress
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:          2
Received prefixes:        2
Suppressed due to damping: 0
Table LDP.inet.0 Bit: 50000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:          1
Received prefixes:        1
Suppressed due to damping: 0
Table OSPF.inet.0 Bit: 60000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:          2
Received prefixes:        2
Suppressed due to damping: 0
Table RIP.inet.0 Bit: 70000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:          2
Received prefixes:        2
Suppressed due to damping: 0
Table STATIC.inet.0 Bit: 80000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:          1
Received prefixes:        1
Suppressed due to damping: 0
Table L2VPN.l2vpn.0 Bit: 90000
RIB State: BGP restart is complete
RIB State: VPN restart in progress
Send state: in sync
Active prefixes:          1
Received prefixes:        1
Suppressed due to damping: 0
Last traffic (seconds): Received 0    Sent 0    Checked 0
Input messages: Total 14    Updates 13    Refreshes 0    Octets 1053
Output messages: Total 3    Updates 0    Refreshes 0    Octets 105
Output Queue[0]: 0 (bgp.l3vpn.0, inet-vpn-unicast)
Output Queue[1]: 0 (bgp.l2vpn.0, inet-vpn-unicast)
Output Queue[2]: 0 (BGP-INET.inet.0, inet-vpn-unicast)
Output Queue[3]: 0 (BGP-L.inet.0, inet-vpn-unicast)
Output Queue[4]: 0 (LDP.inet.0, inet-vpn-unicast)
Output Queue[5]: 0 (OSPF.inet.0, inet-vpn-unicast)
Output Queue[6]: 0 (RIP.inet.0, inet-vpn-unicast)
Output Queue[7]: 0 (STATIC.inet.0, inet-vpn-unicast)
Output Queue[8]: 0 (L2VPN.l2vpn.0, inet-vpn-unicast)

```


show bgp neighbor (Layer 3 VPN)

```

user@host> show bgp neighbor
Peer: 4.4.4.4+179      AS 10045 Local: 5.5.5.5+1214      AS 10045
  Type: Internal      State: Established      Flags: <ImportEval>
  Last State: OpenConfirm  Last Event: RecvKeepAlive
  Last Error: None
  Export: [ match-all ] Import: [ match-all ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
    Rib-group Refresh>
  Address families configured: inet-vpn-unicast
  Local Address: 5.5.5.5 Holdtime: 90 Preference: 170
  Flags for NLRI inet-labeled-unicast: TrafficStatistics
  Traffic Statistics: Options: all File: /var/log/bstat.log
                                size 131072 files 10

  Traffic Statistics Interval: 60
  Number of flaps: 0
  Peer ID: 192.168.1.110      Local ID: 192.168.1.111      Active Holdtime: 90
  Keepalive Interval: 30
  NLRI for restart configured on peer: inet-vpn-unicast
  NLRI advertised by peer: inet-vpn-unicast
  NLRI for this session: inet-vpn-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 120
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-vpn-unicast
  NLRI peer can save forwarding state: inet-vpn-unicast
  NLRI that peer saved forwarding for: inet-vpn-unicast
  NLRI that restart is negotiated for: inet-vpn-unicast
  NLRI of received end-of-rib markers: inet-vpn-unicast
  NLRI of all end-of-rib markers sent: inet-vpn-unicast
  Table bgp.l3vpn.0 Bit: 10000
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: in sync
    Active prefixes:          2
    Received prefixes:        2
    Suppressed due to damping: 0
  Table vpn-green.inet.0 Bit: 20001
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: in sync
    Active prefixes:          2
    Received prefixes:        2
    Suppressed due to damping: 0
  Last traffic (seconds): Received 15      Sent 20      Checked 20
  Input messages: Total 40      Updates 2      Refreshes 0      Octets 856
  Output messages: Total 44      Updates 2      Refreshes 0      Octets 1066
  Output Queue[0]: 0 (bgp.l3vpn.0, inet-vpn-unicast)
  Output Queue[1]: 0 (vpn-green.inet.0, inet-vpn-unicast)
  Trace options: detail packets
  Trace file: /var/log/bgpgr.log size 131072 files 10

```

show bgp neighbor neighbor-address

```

user@host> show bgp neighbor 192.168.1.111
Peer: 10.255.245.12+179 AS 35 Local: 10.255.245.13+2884 AS 35
  Type: Internal      State: Established (route reflector client)Flags: <Sync>
  Last State: OpenConfirm  Last Event: RecvKeepAlive
  Last Error: None

```

```

Options: <Preference LocalAddress HoldTime Cluster AddressFamily Rib-group
Refresh>
Address families configured: inet-vpn-unicast inet-labeled-unicast
Local Address: 10.255.245.13 Holdtime: 90 Preference: 170
Flags for NLRI inet-vpn-unicast: AggregateLabel
Flags for NLRI inet-labeled-unicast: AggregateLabel
Number of flaps: 0
Peer ID: 10.255.245.12    Local ID: 10.255.245.13    Active Holdtime: 90
Keepalive Interval: 30
BFD: disabled
NLRI advertised by peer: inet-vpn-unicast inet-labeled-unicast
NLRI for this session: inet-vpn-unicast inet-labeled-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 300
Stale routes from peer are kept for: 60
Restart time requested by this peer: 300
NLRI that peer supports restart for: inet-unicast inet6-unicast
NLRI that restart is negotiated for: inet-unicast inet6-unicast
NLRI of received end-of-rib markers: inet-unicast inet6-unicast
NLRI of all end-of-rib markers sent: inet-unicast inet6-unicast
Table inet.0 Bit: 10000
  RIB State: restart is complete
  Send state: in sync
  Active prefixes: 4
  Received prefixes: 6
  Suppressed due to damping: 0
Table inet6.0 Bit: 20000
  RIB State: restart is complete
  Send state: in sync
  Active prefixes: 0
  Received prefixes: 2
  Suppressed due to damping: 0
Last traffic (seconds): Received 3    Sent 3    Checked 3
Input messages: Total 9    Updates 6    Refreshes 0    Octets 403
Output messages: Total 7    Updates 3    Refreshes 0    Octets 365
Output Queue[0]: 0 (inet.0, inet-unicast)
Output Queue[1]: 0 (inet6.0, inet6-unicast)
Trace options: detail packets
Trace file: /var/log/bgpr size 131072 files 10

```

show bgp neighbor neighbor-address

```

user@host> show bgp neighbor 192.168.4.222
Peer: 192.168.4.222+4902 AS 65501 Local: 192.168.4.221+179 AS 65500
Type: External State: Established Flags: <Sync>
Last State: OpenConfirm Last Event: RecvKeepAlive
Last Error: Cease
Export: [ export-policy ] Import: [ import-policy ]
Options: <Preference HoldTime AddressFamily PeerAS PrefixLimit Refresh>
Address families configured: inet-unicast inet-multicast
Holdtime: 60000 Preference: 170
Number of flaps: 4
Last flap event: RecvUpdate
Error: 'Cease' Sent: 5 Recv: 0
Peer ID: 10.255.245.6    Local ID: 10.255.245.5    Active Holdtime: 60000
Keepalive Interval: 20000 Peer index: 0
BFD: disabled, down
Local Interface: fxp0.0
NLRI advertised by peer: inet-unicast inet-multicast
NLRI for this session: inet-unicast inet-multicast
Peer supports Refresh capability (2)

```

```

Table inet.0 Bit: 10000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          8
  Received prefixes:       10
  Accepted prefixes:       10
  Suppressed due to damping: 0
  Advertised prefixes:     3
Table inet.2 Bit: 20000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:          0
  Received prefixes:       0
  Accepted prefixes:       0
  Suppressed due to damping: 0
  Advertised prefixes:     0
Last traffic (seconds): Received 357 Sent 357 Checked 357
Input messages: Total 4 Updates 2 Refreshes 0 Octets 211
Output messages: Total 4 Updates 1 Refreshes 0 Octets 147
Output Queue[0]: 0 (inet.0, inet-unicast)
Output Queue[1]: 0 (inet.2, inet-multiicast)
Trace options: all
Trace file: /var/log/bgp size 10485760 files 10

```

show bgp neighbor orf neighbor-address detail

```

user@host > show bgp neighbor orf 192.168.165.56 detail
Peer: 192.168.165.56+179 Type: External
Group: ext1

inet-unicast
  Filter updates rcv:          1 Immediate:          1
  Filter: prefix-based receive
  Received filter entries:
    seq 1: prefix 2.2.2.2/32: minlen 32: maxlen 32: match deny:

inet6-unicast
  Filter updates rcv:          0 Immediate:          1
  Filter: prefix-based receive
  Received filter entries:
    *.*

```

show bgp replication

Syntax	show bgp replication
Release Information	Command introduced in JUNOS Release 8.5. Command introduced in Junos OS Release 11.3 for the QFX Series. Support for logical-system option introduced in Junos OS Release 13.3
Description	Displays the status of BGP state replication between the master and backup Routing Engines on devices that have nonstop active routing configured on them.
Options	logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> <i>show bgp replication logical-system</i>
List of Sample Output	show bgp replication (for Master) on page 339 show bgp replication (for Backup) on page 339
Output Fields	Table 20 on page 338 lists the output fields for the show bgp replication command. Output fields are listed in the approximate order in which they appear.

Table 20: show bgp replication Output Fields

Field Name	Field Description
session state	State of the current internal BGP state replication session, Up or Down, and the duration for which the session has been in the indicated state.
flaps	Total number of flaps that occurred.
protocol state	Current state of the protocol operation, Active, Connect, Idle, and the duration for which the protocol has been in the indicated state.
synchronization state	Synchronization state at the time of executing the command. The states can be: <ul style="list-style-type: none"> Idle Neighbor—Indicates that the neighbor state synchronization is in progress. AckWait—Indicates that the request processing is over. ORF—Indicates that the outbound routing filter synchronization is in progress. RIB—Indicates that the routing table synchronization is in progress. Complete

Table 20: show bgp replication Output Fields (*continued*)

Field Name	Field Description
number of peers waiting	<p>Total number of peers waiting for various messages:</p> <ul style="list-style-type: none"> • AckWait—Number of peers waiting for a connection establishment or completed acknowledgment messages. • SoWait—Number of peers waiting for TCP socket-related operations. • Scheduled—Number of peers being synchronized.
messages sent	<p>Number of various types of messages that have been sent since internal replication session became active:</p> <ul style="list-style-type: none"> • Open—Number of Open messages sent. • Establish—Number of connection establishment acknowledgment messages sent. • Update—Number of update messages sent. • Error—Number of error messages sent. • Complete—Number of connection complete acknowledgment messages sent.
messages received	<p>Total number of messages received:</p> <ul style="list-style-type: none"> • Open—Number of Open messages received. • Request—Number of request messages received: <ul style="list-style-type: none"> • Wildcard—Number of requests received that used wildcards in the target address. • Targeted—Number of requests received that used a specific address. • EstablishAck—Number of connection establishment acknowledgement messages received. • CompleteAck—Number of connection completed acknowledgement messages received.

Sample Output

show bgp replication (for Master)

```

user@host> show bgp replication
Synchronization master:
  Session state: Up, Since: 44:07
  Flaps: 0
  Protocol state: Idle, Since: 14
  Synchronization state: Complete
  Number of peers waiting: AckWait: 0, SoWait: 0, Scheduled: 0
  Messages sent: Open 1, Establish 924, Update 381, Error 60, Complete 114
  Messages received: Open 1, Request 1 wildcard 113 targeted, EstablishAck 924,
  CompleteAck 114

```

show bgp replication (for Backup)

```

user@host> show bgp replication
Synchronization backup:
  State: Established 13 ago
  , Unsync timer: 2

  Unsync entry queue:
    Instance: 0 Neighbor: 30.30.30.1 elapsed: 7
    Instance: 0 Neighbor: 40.40.40.3 elapsed: 7
    Instance: 0 Neighbor: 40.40.40.4 elapsed: 7

```

Instance: 0 Neighbor: 40.40.40.5 elapsed: 7
Instance: 0 Neighbor: 40.40.40.6 elapsed: 7
Instance: 0 Neighbor: 40.40.40.1 elapsed: 7
Instance: 0 Neighbor: 40.40.40.2 elapsed: 7

show bgp summary

Syntax	<pre>show bgp summary <exact-instance <i>instance-name</i>> <group <i>group-name</i>> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switch and QFX Series)	<pre>show bgp summary <exact-instance <i>instance-name</i>> <instance <i>instance-name</i>></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>exact-instance option introduced in Junos OS Release 11.4.</p> <p>group option introduced in Junos OS Release 13.3.</p>
Description	Display BGP summary information.
Options	<p>none—Display BGP summary information for all routing instances.</p> <p>exact-instance <i>instance-name</i>—(Optional) Display information for the specified instance only.</p> <p>group—Display overview of bgp information for a particular group</p> <p>instance <i>instance-name</i>—(Optional) Display information for all routing instances whose name begins with this string (for example, cust1, cust11, and cust111 are all displayed when you run the show bgp summary instance cust1 command). The instance name can be master for the main instance, or any valid configured instance name or its prefix.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	<p>show bgp summary (When a Peer Is Not Established) on page 344</p> <p>show bgp summary (When a Peer Is Established) on page 344</p> <p>show bgp summary (CLNS) on page 344</p> <p>show bgp summary (Layer 2 VPN) on page 345</p> <p>show bgp summary (Layer 3 VPN) on page 345</p> <p>show bgp summary group on page 345</p>
Output Fields	<p>Table 21 on page 342 describes the output fields for the show bgp summary command. Output fields are listed in the approximate order in which they appear.</p>

Table 21: show bgp summary Output Fields

Field Name	Field Description
Groups	Number of BGP groups.
Peers	Number of BGP peers.
Down peers	Number of down BGP peers.
Table	Name of routing table.
Tot Paths	Total number of paths.
Act Paths	Number of active routes.
Suppressed	Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.
History	Number of withdrawn routes stored locally to keep track of damping history.
Damp State	Number of routes with a figure of merit greater than zero, but still active because the value has not reached the threshold at which suppression occurs.
Pending	Routes in process by BGP import policy.
Peer	Address of each BGP peer. Each peer has one line of output.
AS	Peer's AS number.
InPkt	Number of packets received from the peer.
OutPkt	Number of packets sent to the peer.
OutQ	Number of BGP packets that are queued to be transmitted to a particular neighbor. It normally is 0 because the queue usually is emptied quickly.
Flaps	Number of times the BGP session has gone down and then come back up.
Last Up/Down	Last time since the neighbor transitioned to or from the established state.

Table 21: show bgp summary Output Fields (*continued*)

Field Name	Field Description
State #Active /Received/Accepted /Damped	<p>Multipurpose field that displays information about BGP peer sessions. The field's contents depend upon whether a session is established and whether it was established on the main routing device or in a routing instance.</p> <ul style="list-style-type: none"> If a peer is not established, the field shows the state of the peer session: Active, Connect, or Idle. In general, the Idle state is the first stage of a connection. BGP is waiting for a Start event. A session can be idle for other reasons as well. The reason that a session is idle is sometimes displayed. For example: Idle (Removal in progress) or Idle (LicenseFailure). If a BGP session is established on the main routing device, the field shows the number of active, received, accepted, and damped routes that are received from a neighbor and appear in the inet.0 (main) and inet.2 (multicast) routing tables. For example, 8/10/10/2 and 2/4/4/0 indicate the following: <ul style="list-style-type: none"> 8 active routes, 10 received routes, 10 accepted routes, and 2 damped routes from a BGP peer appear in the inet.0 routing table. 2 active routes, 4 received routes, 4 accepted routes, and no damped routes from a BGP peer appear in the inet.2 routing table. If a BGP session is established in a routing instance, the field indicates the established (Establ) state, identifies the specific routing table that receives BGP updates, and shows the number of active, received, and damped routes that are received from a neighbor. For example, Establ VPN-AB.inet.0: 2/4/0 indicates the following: <ul style="list-style-type: none"> The BGP session is established. Routes are received in the VPN-AB.inet.0 routing table. The local routing device has two active routes, four received routes, and no damped routes from a BGP peer. <p>When a BGP session is established, the peers are exchanging update messages.</p>

Sample Output

show bgp summary (When a Peer Is Not Established)

```

user@host> show bgp summary
Groups: 2 Peers: 4 Down peers: 1
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0 6 4 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.0.0.3 65002 86 90 0 2 42:54 0/0/0

0/0/0
10.0.0.4 65002 90 91 0 1 42:54 0/2/0

0/0/0
10.0.0.6 65002 87 90 0 3 3 Active
10.1.12.1 65001 89 89 0 1 42:54 4/4/0

0/0/0

```

show bgp summary (When a Peer Is Established)

```

user@host> show bgp summary
Groups: 1 Peers: 3 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0 6 4 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.0.0.2 65002 88675 88652 0 2 42:38 2/4/0

0/0/0
10.0.0.3 65002 54528 54532 0 1 2w4d22h 0/0/0

0/0/0
10.0.0.4 65002 51597 51584 0 0 2w3d22h 2/2/0

0/0/0

user@host> show bgp summary logical-system R3
Groups: 2 Peers: 2 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
bgp.13vpn.0 2 2 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
1.1.1.2 2 204 206 0 0 1:30:59
Establ
  bgp.13vpn.0: 2/2/2/0
  red.inet.0: 2/2/2/0
10.1.1.10 3 206 207 0 0 1:31:36
Establ
  red.inet.0: 2/2/2/0

```

show bgp summary (CLNS)

```

user@host> show bgp summary
Groups: 1 Peers: 1 Down peers: 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.245.245.1 200 1735 1737 0 0 14:26:12 Establ

```

```

bgp.isovpn.0: 3/3/0
aaaa.iso.0: 3/3/0

```

show bgp summary (Layer 2 VPN)

```

user@host> show bgp summary
Groups: 1 Peers: 5 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
bgp.l2vpn.0 1 1 0 0 0 0
inet.0 0 0 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.255.245.35 65299 72 74 0 1 19:00 Establ
bgp.l2vpn.0: 1/1/0
frame-vpn.l2vpn.0: 1/1/0
10.255.245.36 65299 2164 2423 0 4 19:50 Establ
bgp.l2vpn.0: 0/0/0
frame-vpn.l2vpn.0: 0/0/0
10.255.245.37 65299 36 37 0 4 17:07 Establ
inet.0: 0/0/0
10.255.245.39 65299 138 168 0 6 53:48 Establ
bgp.l2vpn.0: 0/0/0
frame-vpn.l2vpn.0: 0/0/0
10.255.245.69 65299 134 140 0 6 53:42 Establ
inet.0: 0/0/0

```

show bgp summary (Layer 3 VPN)

```

user@host> show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
bgp.l3vpn.0 2 2 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.39.1.5 2 21 22 0 0 6:26 Establ
VPN-AB.inet.0: 1/1/0
10.255.71.15 1 19 21 0 0 6:17 Establ
bgp.l3vpn.0: 2/2/0
VPN-A.inet.0: 1/1/0
VPN-AB.inet.0: 2/2/0
VPN-B.inet.0: 1/1/0

```

show bgp summary group

```

user@host> show bgp summary group Group2
Groups: 3 Peers: 3 Down peers: 3
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0 0 0 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.0.0.1 56 0 0 0 0 51
Idle

user@host> show bgp summary logical-system R3 group toR4
Groups: 2 Peers: 2 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
bgp.l3vpn.0 2 2 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Accepted/Damped...
10.1.1.10 3 207 207 0 0 1:31:40

```

```
Estab1  
red.inet.0: 2/2/2/0
```

Sample Output

show policy damping

Syntax	show policy damping <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch and QFX Series)	show policy damping
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display information about BGP route flap damping parameters.
Options	<p>none—Display information about BGP route flap damping parameters.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Additional Information	In the output from this command, figure-of-merit values correlate with the probability of future instability of a routing device. Routes with higher figure-of-merit values are suppressed for longer periods of time. The figure-of-merit value decays exponentially over time. A figure-of-merit value of zero is assigned to each new route. The value is increased each time the route is withdrawn or readvertised, or when one of its path attributes changes.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • “Configuring BGP Flap Damping Parameters” in the <i>Routing Policy Feature Guide for Routing Devices</i> • clear bgp damping on page 308 • show route damping on page 802
List of Sample Output	show policy damping on page 348
Output Fields	Table 22 on page 347 describes the output fields for the show policy damping command. Output fields are listed in the approximate order in which they appear.

Table 22: show policy damping Output Fields

Field Name	Field Description
Halflife	Decay half-life, in minutes. The value represents the period during which the accumulated figure-of-merit value is reduced by half if the route remains stable. If a route has flapped, but then becomes stable, the figure-of-merit value for the route decays exponentially. For example, for a route with a figure-of-merit value of 1500, if no incidents occur, its figure-of-merit value is reduced to 750 after 15 minutes and to 375 after another 15 minutes.

Table 22: show policy damping Output Fields (*continued*)

Field Name	Field Description
Reuse merit	Figure-of-merit value below which a suppressed route can be used again. A suppressed route becomes reusable when its figure-of-merit value decays to a value below a reuse threshold, and the route once again is considered usable and can be installed in the forwarding table and exported from the routing table.
Suppress/cutoff merit	Figure-of-merit value above which a route is suppressed for use or inclusion in advertisements. When a route's figure-of-merit value reaches a particular level, called the cutoff or suppression threshold, the route is suppressed. When a route is suppressed, the routing table no longer installs the route into the forwarding table and no longer exports this route to any of the routing protocols.
Maximum suppress time	Maximum hold-down time, in minutes. The value represents the maximum time that a route can be suppressed no matter how unstable it has been before this period of stability.
Computed values	<ul style="list-style-type: none"> • Merit ceiling—Maximum merit that a flapping route can collect. • Maximum decay—Maximum decay half-life, in minutes.

Sample Output

show policy damping

```

user@host> show policy damping
Default damping information:
  Halflife: 15 minutes
  Reuse merit: 750 Suppress/cutoff merit: 3000
  Maximum suppress time: 60 minutes
  Computed values:
    Merit ceiling: 12110
    Maximum decay: 6193
Damping information for "standard-damping":
  Halflife: 10 minutes
  Reuse merit: 4000 Suppress/cutoff merit: 8000
  Maximum suppress time: 30 minutes
  Computed values:
    Merit ceiling: 32120
    Maximum decay: 12453

```

CHAPTER 9

ES-IS Operational Mode Commands

- clear esis adjacency
- clear esis statistics
- show esis adjacency
- show esis interface
- show esis statistics

clear esis adjacency

Syntax	clear esis adjacency <instance <i>instance-name</i> > <interface <i>interface-name</i> > < <i>neighbor</i> >
Release Information	Command introduced before Junos OS Release 7.4.
Description	Clear End System-to-Intermediate System (ES-IS) adjacencies.
Options	none —Clear all ES-IS adjacencies. instance <i>instance-name</i> —(Optional) Clear adjacencies for the specified routing instance only. interface <i>interface-name</i> —(Optional) Clear adjacencies for the specified interface only. <i>neighbor</i> —(Optional) Clear adjacencies for the specified neighbor only.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show esis adjacency on page 352
List of Sample Output	clear esis adjacency on page 350
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear esis adjacency

```
user@host> clear esis adjacency
```


clear esis statistics

Syntax	clear esis statistics <instance <i>instance-name</i> >
Release Information	Command introduced before Junos OS Release 7.4.
Description	Clear End System-to-Intermediate System (ES-IS) packet statistics.
Options	none —Clear ES-IS packet statistics for all routing instances. instance <i>instance-name</i> —(Optional) Clear ES-IS packet statistics for the specified routing instance only.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show esis statistics on page 356
List of Sample Output	clear esis statistics on page 351
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear esis statistics

```
user@host> clear esis statistics
```

show esis adjacency

Syntax	show esis adjacency <brief detail extensive> <esis-neighbor-id> <instance <i>instance-name</i> > <interface <i>interface-name</i> >
Release Information	Command introduced before Junos OS Release 7.4.
Description	(J Series and PTX Series routers only) Display End System-to-Intermediate System (ES-IS) adjacencies.
Options	<p>none—(Same as brief) Display all ES-IS adjacencies.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>esis-neighbor-id—(Optional) Display adjacencies for the specified neighbor's network service access point (NSAP) only.</p> <p>instance <i>instance-name</i>—(Optional) Display adjacencies for the specified routing instance only.</p> <p>interface <i>interface-name</i>—(Optional) Display adjacencies for the specified interface only.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> clear esis adjacency on page 350
List of Sample Output	show esis adjacency on page 353 show esis adjacency brief on page 353 show esis adjacency detail on page 353 show esis adjacency extensive on page 353
Output Fields	Table 23 on page 352 describes the output fields for the show esis adjacency command. Output fields are listed in the approximate order in which they appear.

Table 23: show esis adjacency Output Fields

Field Name	Field Description	Level of Output
Nbr Type	Type of network service access point (NSAP) of this neighbor.	brief none
NSAP/NET	NSAP of this neighbor.	All levels
Type	Type of NSAP of this neighbor.	detail extensive
Hold (secs)	Holdtime interval advertised by this neighbor.	brief none
Interface	Interface through which the neighbor is reachable.	All levels

Table 23: show esis adjacency Output Fields (*continued*)

Field Name	Field Description	Level of Output
Advertised holdtime	Holdtime interval advertised by this neighbor.	detail extensive
Expires in	How long until the adjacency expires, in seconds.	detail extensive
SNPA	Subnetwork point of attachment (MAC address of the neighbor).	detail extensive
Transition log	List of recent transitions. <ul style="list-style-type: none"> • When—Time of advertisement from this neighbor. • State—State of the adjacency: Up, Down, New, One-way, Initializing, or Rejected. • Event—Event causing the state. • Down reason—Reason the adjacency is down. 	extensive

Sample Output

show esis adjacency

```

user@host> show esis adjacency
Nbr      NSAP/NET                      Hold  Interface
Type                                           (secs)
IS       47.0005.80ff.f800.0000.0108.0001.0102.5501.6008    135 fe-0/0/0.0

```

show esis adjacency brief

The output for the **show esis adjacency brief** command is identical to that for the **show esis adjacency** command. For sample output, see [show esis adjacency on page 353](#).

show esis adjacency detail

```

user@host> show esis adjacency detail
NSAP/NET: 47.0005.80ff.f800.0000.0108.0001.0102.5501.6008, Type: IS
Interface: fe-0/0/0.0, Advertised hold time: 180 secs, Expires in: 173 secs
SNPA: 0:5:85:c1:73:71

```

show esis adjacency extensive

```

user@host> show esis adjacency extensive
NSAP/NET: 47.0005.80ff.f800.0000.0108.0001.0102.5501.6008, Type: IS
Interface: fe-0/0/0.0, Advertised hold time: 180 secs, Expires in: 167 secs
SNPA: 0:5:85:c1:73:71
Transition log:
When           State      Event           Down reason
Sun Nov 26 22:07:35  Up        Received ISH

```

show esis interface

Syntax	show esis interface <brief detail extensive> <instance <i>instance-name</i> > <interface <i>interface-name</i> >
Release Information	Command introduced before Junos OS Release 7.4.
Description	(J Series and PTX Series routers only) Display End System-to-Intermediate System (ES-IS) interface information.
Options	<p>none—(Same as brief) Display information for all configured ES-IS interfaces.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display configured interfaces for the specified routing instance only.</p> <p>interface <i>interface-name</i>—(Optional) Display information about the specified interface only.</p>
Required Privilege Level	view
List of Sample Output	show esis interface on page 355 show esis interface brief on page 355 show esis interface detail on page 355 show esis interface extensive on page 355
Output Fields	Table 24 on page 354 describes the output fields for the show esis interface command. Output fields are listed in the approximate order in which they appear.

Table 24: show esis interface Output Fields

Field Name	Field Description	Level of Output
Interface	Interface through which the adjacency is made.	All levels
Receives	Types of hello messages that are received.	All levels
Sends	Types of hello messages that are sent.	All levels
Hello interval	Interface's hello interval, in seconds.	All levels
Adjacencies or Num Adj	Number of adjacencies established on this interface.	All levels
Holdtime	Interface's hold time, in seconds.	detail extensive
State	Internal implementation information.	detail extensive

Table 24: show esis interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
End system configuration timer	Time, in seconds, for the end system to configure itself for ES-IS.	detail extensive
Interface index	Index value.	detail extensive
NET used in hello	Network entity title used in hello messages.	detail extensive

Sample Output

show esis interface

```

user@host> show esis interface
Interface           Receives    Sends    Hello Interval    Num Adj
fe-0/0/0.0          ISH         ISH         60.00             1
lo0.0               ISH         -          60.00             0

```

show esis interface brief

The output for the **show esis interface brief** command is identical to that for the **show esis interface** command. For sample output, see [show esis interface on page 355](#).

show esis interface detail

```

user@host> show esis interface detail
Interface: fe-0/0/0.0
  Receives: ISH, Sends: ISH, Hello interval: 60.00
  Adjacencies: 1, Holdtime: 180, End system configuration timer: 180
  Interface index: 68, State: 0x2
  NET used in hello: 47.0005.80ff.f800.0000.0108.0001.0102.5501.6007

Interface: lo0.0
  Receives: ISH, Sends: - , Hello interval: 60.00
  Adjacencies: 0, Holdtime: 180, End system configuration timer: 180
  Interface index: 64, State: 0x2
  NET used in hello: 47.0005.80ff.f800.0000.0108.0001.0102.5501.6007

```

show esis interface extensive

The output for the **show esis interface extensive** command is identical to that for the **show esis interface detail** command. For sample output, see [show esis interface detail on page 355](#).

show esis statistics

Syntax	<code>show esis statistics</code> <code><instance <i>instance-name</i>></code>
Release Information	Command introduced before Junos OS Release 7.4.
Description	(J Series and PTX Series routers only) Display End System-to-Intermediate System (ES-IS) packet statistics.
Options	<p>none—Display ES-IS packet statistics for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Display ES-IS statistics for the specified routing instance only.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> clear esis statistics on page 351
List of Sample Output	show esis statistics on page 356
Output Fields	Table 25 on page 356 describes the output fields for the show esis statistics command. Output fields are listed in the approximate order in which they appear.

Table 25: show esis statistics Output Fields

Field Name	Field Description
PDU type	Protocol data unit type.
Received	Number of PDUs received since IS-IS started or since the statistics were set to zero.
Processed	Number of PDUs received less the number dropped.
Drops	Number of PDUs dropped.
Sent	Number of PDUs transmitted since IS-IS started or since the statistics were set to zero.
Total packets received/sent	Total number of PDUs received and transmitted since IS-IS started or since the statistics were set to zero.

Sample Output

show esis statistics

```

user@host> show esis statistics
PDU type  Received  Processed  Drops  Sent
ESH              3          3      0      8

```

ISH	11	10	1	4
RD	0	0	0	0
Unknown	0	0	0	0
Totals	14	13	1	12

Total packets received: 14 sent: 0

CHAPTER 10

IP Multicast Operational Mode Commands

- clear amt statistics
- clear amt tunnel
- clear igmp membership
- clear igmp snooping membership
- clear igmp snooping statistics
- clear igmp statistics
- clear mld membership
- clear mld statistics
- clear msdp cache
- clear msdp statistics
- clear multicast bandwidth-admission
- clear multicast forwarding-cache
- clear multicast scope
- clear multicast sessions
- clear multicast snooping statistics
- clear multicast statistics
- clear pgm negative-acknowledgments
- clear pgm source-path-messages
- clear pgm statistics
- clear pim join
- clear pim join-distribution
- clear pim register
- clear pim snooping join
- clear pim snooping statistics
- clear pim statistics
- request pim multicast-tunnel rebalance

- `show amt statistics`
- `show amt summary`
- `show amt tunnel`
- `show dvmrp interfaces`
- `show dvmrp neighbors`
- `show dvmrp prefix`
- `show dvmrp prunes`
- `show igmp group`
- `show igmp interface`
- `show igmp snooping interface`
- `show igmp snooping membership`
- `show igmp snooping statistics`
- `show igmp statistics`
- `show mld group`
- `show mld interface`
- `show mld statistics`
- `show msdp`
- `show msdp source`
- `show msdp source-active`
- `show msdp statistics`
- `show multicast backup-pe-groups`
- `show multicast flow-map`
- `show multicast forwarding-cache statistics`
- `show multicast interface`
- `show multicast mrinfo`
- `show multicast next-hops`
- `show multicast pim-to-igmp-proxy`
- `show multicast pim-to-mld-proxy`
- `show multicast route`
- `show multicast rpf`
- `show multicast scope`
- `show multicast snooping next-hops`
- `show multicast sessions`
- `show multicast snooping route`
- `show multicast statistics`
- `show multicast usage`
- `show pgm negative-acknowledgments`

- `show pgm source-path-messages`
- `show pgm statistics`
- `show pim bootstrap`
- `show pim interfaces`
- `show pim join`
- `show pim mdt`
- `show pim mdt data-mdt-joins`
- `show pim mdt data-mdt-limit`
- `show pim neighbors`
- `show pim rps`
- `show pim snooping interfaces`
- `show pim snooping join`
- `show pim snooping neighbors`
- `show pim snooping statistics`
- `show pim source`
- `show pim statistics`
- `show sap listen`
- `test msdp`

clear amt statistics

Syntax	clear amt statistics <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced in JUNOS Release 10.2.
Description	Clear Automatic Multicast Tunneling (AMT) statistics.
Options	none —Clear the multicast statistics for all AMT tunnel interfaces. instance <i>instance-name</i> —(Optional) Clear AMT multicast statistics for the specified instance. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show amt statistics on page 398
List of Sample Output	clear amt statistics on page 362
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear amt statistics

```
user@host> clear amt statistics
```

clear amt tunnel

Syntax	<pre>clear amt tunnel <gateway <i>gateway-ip-addr</i>> <port <i>port-number</i>> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <statistics> <tunnel-interface <i>interface-name</i>></pre>
Release Information	Command introduced in JUNOS Release 10.2.
Description	Clear the Automatic Multicast Tunneling (AMT) multicast state. Optionally, clear AMT protocol statistics.
Options	<p>none—Clear multicast state for all AMT tunnel interfaces.</p> <p>gateway <i>gateway-ip-addr</i> port <i>port-number</i>—(Optional) Clear the AMT multicast state for the specified gateway address. If no port is specified, clear the AMT multicast state for all AMT gateways with the given IP address.</p> <p>instance <i>instance-name</i>—(Optional) Clear the AMT multicast state for the specified instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>statistics—(Optional) Clear multicast statistics for all AMT tunnels or for specified tunnels.</p> <p>tunnel-interface <i>interface-name</i>—(Optional) Clear the AMT multicast state for the specified AMT tunnel interface.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show amt tunnel on page 403
List of Sample Output	clear amt tunnel on page 363 clear amt tunnel statistics gateway-address on page 363
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear amt tunnel

```
user@host> clear amt tunnel
```

clear amt tunnel statistics gateway-address

```
user@host> clear amt tunnel statistics gateway-address 100.31.1.21 port 4000
```

clear igmp membership

Syntax	clear igmp membership <group <i>address-range</i> > <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch and the QFX Series)	clear igmp membership <group <i>address-range</i> > <interface <i>interface-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Clear Internet Group Management Protocol (IGMP) group members.
Options	none —Clear all IGMP members on all interfaces and for all address ranges. group <i>address-range</i> —(Optional) Clear all IGMP members that are in a particular address range. An example of a range is 224.2/16 . If you omit the destination prefix length, the default is /32 . interface <i>interface-name</i> —(Optional) Clear all IGMP group members on an interface. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show igmp group on page 414• show igmp interface on page 418
List of Sample Output	clear igmp membership on page 364 clear igmp membership interface on page 365 clear igmp membership group on page 365
Output Fields	See show igmp group for an explanation of output fields.

Sample Output

clear igmp membership

The following sample output displays IGMP group information before and after the **clear igmp membership** command is entered:

```
user@host> show igmp group
Interface      Group           Last Reported  Timeout
so-0/0/0       224.2.127.253   10.1.128.1     186
so-0/0/0       224.2.127.254   10.1.128.1     186
so-0/0/0       239.255.255.255 10.1.128.1     187
```

so-0/0/0	224.1.127.255	10.1.128.1	188
local	224.0.0.6	(null)	0
local	224.0.0.5	(null)	0
local	224.2.127.254	(null)	0
local	239.255.255.255	(null)	0
local	224.0.0.2	(null)	0
local	224.0.0.13	(null)	0

```
user@host> clear igmp membership
Clearing Group Membership Info for so-0/0/0
Clearing Group Membership Info for so-1/0/0
Clearing Group Membership Info for so-2/0/0
```

```
user@host> show igmp group
```

Interface	Group	Last Reported	Timeout
local	224.0.0.6	(null)	0
local	224.0.0.5	(null)	0
local	224.2.127.254	(null)	0
local	239.255.255.255	(null)	0
local	224.0.0.2	(null)	0
local	224.0.0.13	(null)	0

clear igmp membership interface

The following sample output displays IGMP group information before and after the **clear igmp membership interface** command is issued:

```
user@host> show igmp group
```

Interface	Group	Last Reported	Timeout
so-0/0/0	224.2.127.253	10.1.128.1	210
so-0/0/0	239.255.255.255	10.1.128.1	210
so-0/0/0	224.1.127.255	10.1.128.1	215
so-0/0/0	224.2.127.254	10.1.128.1	216
local	224.0.0.6	(null)	0
local	224.0.0.5	(null)	0
local	224.2.127.254	(null)	0
local	239.255.255.255	(null)	0
local	224.0.0.2	(null)	0
local	224.0.0.13	(null)	0

```
user@host> clear igmp membership interface so-0/0/0
Clearing Group Membership Info for so-0/0/0
```

```
user@host> show igmp group
```

Interface	Group	Last Reported	Timeout
local	224.0.0.6	(null)	0
local	224.0.0.5	(null)	0
local	224.2.127.254	(null)	0
local	239.255.255.255	(null)	0
local	224.0.0.2	(null)	0
local	224.0.0.13	(null)	0

clear igmp membership group

The following sample output displays IGMP group information before and after the **clear igmp membership group** command is entered:

```
user@host> show igmp group
```

Interface	Group	Last Reported	Timeout
so-0/0/0	224.2.127.253	10.1.128.1	210
so-0/0/0	239.255.255.255	10.1.128.1	210
so-0/0/0	224.1.127.255	10.1.128.1	215
so-0/0/0	224.2.127.254	10.1.128.1	216
local	224.0.0.6	(null)	0
local	224.0.0.5	(null)	0
local	224.2.127.254	(null)	0
local	239.255.255.255	(null)	0
local	224.0.0.2	(null)	0
local	224.0.0.13	(null)	0

```
user@host> clear igmp membership group 239.225/16
Clearing Group Membership Range 239.225.0.0/16 on so-0/0/0
Clearing Group Membership Range 239.225.0.0/16 on so-1/0/0
Clearing Group Membership Range 239.225.0.0/16 on so-2/0/0
```

```
user@host> show igmp group
```

Interface	Group	Last Reported	Timeout
so-0/0/0	224.1.127.255	10.1.128.1	231
so-0/0/0	224.2.127.254	10.1.128.1	233
so-0/0/0	224.2.127.253	10.1.128.1	236
local	224.0.0.6	(null)	0
local	224.0.0.5	(null)	0
local	224.2.127.254	(null)	0
local	239.255.255.255	(null)	0
local	224.0.0.2	(null)	0
local	224.0.0.13	(null)	0

clear igmp snooping membership

Syntax	clear igmp snooping membership <group source <i>address</i> > <instance <i>instance-name</i> > <interface <i>interface-name</i> > <learning-domain <i>learning-domain-name</i> > <vlan-id <i>vlan-identifier</i> >
Release Information	Command introduced in Junos OS Release 8.5.
Description	Clear IP IGMP snooping membership information.
Options	<p>none—Clear IGMP snooping membership for all supported address families on all interfaces.</p> <p>group source <i>address</i>—(Optional) Clear IGMP snooping membership for the specified multicast group or source address.</p> <p>instance <i>instance-name</i>—(Optional) Clear IGMP snooping membership for the specified instance.</p> <p>interface <i>interface-name</i>—(Optional) Clear IGMP snooping membership on a specific interface.</p> <p>learning-domain <i>learning-domain-name</i>—(Optional) Perform this operation on all learning domains or on a particular learning domain.</p> <p>vlan-id <i>vlan-identifier</i>—(Optional) Perform this operation on a particular VLAN.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show igmp snooping membership on page 425
List of Sample Output	clear igmp snooping membership on page 367
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear igmp snooping membership

```
user@host> clear igmp snooping membership
```

clear igmp snooping statistics

Syntax	clear igmp snooping statistics <instance <i>instance-name</i> > <interface <i>interface-name</i> > <learning-domain (all <i>learning-domain-name</i>)> <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced in Junos OS Release 8.5.
Description	Clear IP IGMP snooping statistics.
Options	<p>none—Clear IGMP snooping statistics for all supported address families on all interfaces.</p> <p>instance <i>instance-name</i>—(Optional) Clear IGMP snooping statistics for the specified instance.</p> <p>interface <i>interface-name</i>—(Optional) Clear IGMP snooping statistics on a specific interface.</p> <p>learning-domain (all <i>learning-domain-name</i>)—(Optional) Perform this operation on all learning domains or on a particular learning domain.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show igmp snooping statistics on page 429
List of Sample Output	clear igmp snooping statistics on page 368
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear igmp snooping statistics

```
user@host> clear igmp snooping statistics
```

clear igmp statistics

Syntax	clear igmp statistics <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	clear igmp statistics <interface <i>interface-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Clear Internet Group Management Protocol (IGMP) statistics.
Options	<p>none—Clear IGMP statistics on all interfaces.</p> <p>interface <i>interface-name</i>—(Optional) Clear IGMP statistics for the specified interface only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show igmp statistics on page 432
List of Sample Output	clear igmp statistics on page 369
Output Fields	See show igmp statistics for an explanation of output fields.

Sample Output

clear igmp statistics

The following sample output displays IGMP statistics information before and after the **clear igmp statistics** command is entered:

```

user@host> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type      Received      Sent  Rx errors
Membership Query        8883         459      0
V1 Membership Report     0            0        0
DVMRP                   19784       35476      0
PIM V1                  18310         0        0
Cisco Trace              0            0        0
V2 Membership Report     0            0        0
Group Leave              0            0        0
Mtrace Response          0            0        0
Mtrace Request           0            0        0
Domain Wide Report       0            0        0
V3 Membership Report     0            0        0

```

Other Unknown types	0
IGMP v3 unsupported type	0
IGMP v3 source required for SSM	0
IGMP v3 mode not applicable for SSM	0

IGMP Global Statistics	
Bad Length	0
Bad Checksum	0
Bad Receive If	0
Rx non-local	1227

user@host> clear igmp statistics

user@host> show igmp statistics

IGMP packet statistics for all interfaces

IGMP Message type	Received	Sent	Rx errors
Membership Query	0	0	0
V1 Membership Report	0	0	0
DVMRP	0	0	0
PIM V1	0	0	0
Cisco Trace	0	0	0
V2 Membership Report	0	0	0
Group Leave	0	0	0
Mtrace Response	0	0	0
Mtrace Request	0	0	0
Domain Wide Report	0	0	0
V3 Membership Report	0	0	0
Other Unknown types			0
IGMP v3 unsupported type			0
IGMP v3 source required for SSM			0
IGMP v3 mode not applicable for SSM			0
IGMP Global Statistics			
Bad Length	0		
Bad Checksum	0		
Bad Receive If	0		
Rx non-local	0		

clear mld membership

Syntax	clear mld membership <group <i>group-name</i> > <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Clear Multicast Listener Discovery (MLD) group membership.
Options	<p>none—Clear all MLD memberships.</p> <p>group <i>group-name</i>—(Optional) Clear MLD membership for the specified group.</p> <p>interface <i>interface-name</i>—(Optional) Clear MLD group membership for the specified interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show mld group on page 435
List of Sample Output	clear mld membership on page 371
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear mld membership

```
user@host> clear mld membership
```

clear mld statistics

Syntax	<code>clear mld statistics</code> <code><interface <i>interface-name</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Clear Multicast Listener Discovery (MLD) statistics.
Options	none —(Same as logical-system all) Clear MLD statistics for all interfaces. interface <i>interface-name</i> —(Optional) Clear MLD statistics for the specified interface. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show mld statistics on page 443
List of Sample Output	clear mld statistics on page 372
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear mld statistics

```
user@host> clear mld statistics
```

clear msdp cache

Syntax	clear msdp cache <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <peer <i>peer-address</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Clear the entries in the Multicast Source Discovery Protocol (MSDP) source-active cache.
Options	<p>none—Clear entries in the MSDP source-active cache for all instances, logical systems, and peers.</p> <p>instance <i>instance-name</i>—(Optional) Clear entries for a specific MSDP instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>peer <i>peer-address</i>—(Optional) Clear the MSDP source-active cache entries learned from a specific peer.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show msdp source-active on page 450
List of Sample Output	clear msdp cache on page 373
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear msdp cache

```
user@host> clear msdp cache
```

clear msdp statistics

Syntax	<code>clear msdp statistics</code> <code><instance <i>instance-name</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code> <code><peer <i>peer-address</i>></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Clear Multicast Source Discovery Protocol (MSDP) peer statistics.
Options	none —Clear MSDP statistics for all peers. instance <i>instance-name</i> —(Optional) Clear statistics for the specified instance. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. peer <i>peer-address</i> —(Optional) Clear the statistics for the specified peer.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show msdp statistics on page 453
List of Sample Output	clear msdp statistics on page 374
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear msdp statistics

```
user@host> clear msdp statistics
```


clear multicast bandwidth-admission

Syntax	<pre>clear multicast bandwidth-admission <group <i>group-address</i>> <inet inet6> <instance <i>instance-name</i>> <interface <i>interface-name</i>> <source <i>source-address</i>></pre>
Release Information	<p>Command introduced in Junos OS Release 8.3.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Reapply IP multicast bandwidth admissions.
Options	<p>none—Reapply multicast bandwidth admissions for all IPv4 forwarding entries in the master routing instance.</p> <p>group <i>group-address</i>—(Optional) Reapply multicast bandwidth admissions for the specified group.</p> <p>inet—(Optional) Reapply multicast bandwidth admission settings for IPv4 flows.</p> <p>inet6—(Optional) Reapply multicast bandwidth admission settings for IPv6 flows.</p> <p>instance <i>instance-name</i>—(Optional) Reapply multicast bandwidth admission settings for the specified instance. If you do not specify an instance, the command applies to the master routing instance.</p> <p>interface <i>interface-name</i>—(Optional) Examines the corresponding outbound interface in the relevant entries and acts as follows:</p> <ul style="list-style-type: none"> • If the interface is congested, and it was admitted previously, it is removed. • If the interface was rejected previously, the clear multicast bandwidth-admission command enables the interface to be admitted as long as enough bandwidth exists on the interface. • If you do not specify an interface, issuing the clear multicast bandwidth-admission command readmits any previously rejected interface for the relevant entries as long as enough bandwidth exists on the interface. <p>To manually reject previously admitted outbound interfaces, you must specify the interface.</p> <p>source <i>source-address</i>—(Optional) Use with the group option to reapply multicast bandwidth admission settings for the specified (source, group) entry.</p>
Required Privilege Level	clear

Related Documentation • [show multicast interface on page 463](#)

List of Sample Output [clear multicast bandwidth-admission on page 376](#)

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

Sample Output

[clear multicast bandwidth-admission](#)

```
user@host> clear multicast bandwidth-admission
```

clear multicast forwarding-cache

Syntax	clear multicast forwarding-cache <inet inet6> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced in Junos OS Release 12.2.
Description	Clear IP multicast forwarding cache entries. This command is not supported for next-generation multiprotocol BGP multicast VPNs (MVPNs).
Options	none —(Same as logical-system all) Clear multicast forwarding cache entries. inet —(Optional) Clear multicast forwarding cache entries for IPv4 family addresses. inet6 —(Optional) Clear multicast forwarding cache entries for IPv6 family addresses. instance <i>instance-name</i> —(Optional) Clear multicast forwarding cache entries on a specific routing instance. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show multicast forwarding-cache statistics on page 461
List of Sample Output	clear multicast forwarding-cache on page 377
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear multicast forwarding-cache

```
user@host> clear multicast forwarding-cache
```

clear multicast scope

Syntax	clear multicast scope <inet inet6> <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch and the QFX Series)	clear multicast scope <inet inet6> <interface <i>interface-name</i> >
Release Information	Command introduced in Junos OS Release 7.6. Command introduced in Junos OS Release 9.0 for EX Series switches. inet6 option introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Clear IP multicast scope statistics.
Options	none —(Same as logical-system all) Clear multicast scope statistics. inet —(Optional) Clear multicast scope statistics for IPv4 family addresses. inet6 —(Optional) Clear multicast scope statistics for IPv6 family addresses. interface <i>interface-name</i> —(Optional) Clear multicast scope statistics on a specific interface. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show multicast scope on page 485
List of Sample Output	clear multicast scope on page 378
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear multicast scope

```
user@host> clear multicast scope
```

clear multicast sessions

Syntax	clear multicast sessions <logical-system (all <i>logical-system-name</i>)> < <i>regular-expression</i> >
Syntax (EX Series Switch and the QFX Series)	clear multicast sessions < <i>regular-expression</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Clear IP multicast sessions.
Options	<p>none—(Same as logical-system all) Clear multicast sessions.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>regular-expression</i>—(Optional) Clear only multicast sessions that contain the specified regular expression.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show multicast sessions on page 489
List of Sample Output	clear multicast sessions on page 379
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear multicast sessions

```
user@host> clear multicast sessions
```

clear multicast snooping statistics

Syntax	clear multicast snooping statistics <instance <i>instance-name</i> > <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced in Junos OS Release 8.5.
Description	Clear IP multicast snooping statistics.
Options	<p>none—Clear multicast snooping statistics for all supported address families on all interfaces.</p> <p>instance <i>instance-name</i>—(Optional) Clear multicast snooping statistics for the specified instance.</p> <p>interface <i>interface-name</i>—(Optional) Clear multicast snooping statistics on a specific interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	clear
List of Sample Output	clear multicast snooping statistics on page 380
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear multicast snooping statistics

```
user@host> clear multicast snooping statistics
```

clear multicast statistics

Syntax	clear multicast statistics <inet inet6> <instance <i>instance-name</i> > <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch and the QFX Series)	clear multicast statistics <inet inet6> <instance <i>instance-name</i> > <interface <i>interface-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Clear IP multicast statistics.
Options	<p>none—Clear multicast statistics for all supported address families on all interfaces.</p> <p>inet—(Optional) Clear multicast statistics for IPv4 family addresses.</p> <p>inet6—(Optional) Clear multicast statistics for IPv6 family addresses.</p> <p>instance <i>instance-name</i>—(Optional) Clear multicast statistics for the specified instance.</p> <p>interface <i>interface-name</i>—(Optional) Clear multicast statistics on a specific interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show multicast statistics on page 495
List of Sample Output	clear multicast statistics on page 381
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear multicast statistics

```
user@host> clear multicast statistics
```

clear pgm negative-acknowledgments

Syntax	clear pgm negative-acknowledgments
Release Information	Command introduced before Junos OS Release 7.4.
Description	Clear the Pragmatic General Multicast (PGM) negative acknowledgment (NAK) state received.
Options	This command has no options.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show pgm negative-acknowledgments on page 501
List of Sample Output	clear pgm negative-acknowledgments on page 382
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear pgm negative-
acknowledgments

```
user@host> clear pgm negative-acknowledgments
```


clear pgm source-path-messages

Syntax	clear pgm source-path-messages
Release Information	Command introduced before Junos OS Release 7.4.
Description	Clear Pragmatic General Multicast (PGM) source-path messages.
Options	This command has no options.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show pgm source-path-messages on page 503
List of Sample Output	clear pgm source-path-messages on page 383
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear pgm source-path-messages

```
user@host> clear pgm source-path-messages
```

clear pgm statistics

Syntax	clear pgm statistics
Release Information	Command introduced before Junos OS Release 7.4.
Description	Clear Pragmatic General Multicast (PGM) statistics.
Options	This command has no options.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show pgm statistics on page 504
List of Sample Output	clear pgm statistics on page 384
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear pgm statistics

```
user@host> clear pgm statistics
```

clear pim join

Syntax	<pre>clear pim join <group-address> <bidirectional dense sparse> <exact> <inet inet6> <instance instance-name> <logical-system (all logical-system-name)> <rp ip-address/prefix source ip-address/prefix> <sg star-g></pre>
Syntax (EX Series Switch and the QFX Series)	<pre>clear pim join <group-address> <dense sparse> <exact> <inet inet6> <instance instance-name> <rp ip-address/prefix source ip-address/prefix> <sg star-g></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Multiple new filter options introduced in Junos OS Release 13.2.</p>
Description	Clear the Protocol Independent Multicast (PIM) join and prune states.
Options	<p>none—Clear the PIM join and prune states for all groups, family addresses, and instances.</p> <p>group-address—(Optional) Clear the PIM join and prune states for a group address.</p> <p>bidirectional dense sparse—(Optional) Clear PIM bidirectional mode, dense mode, or sparse and source-specific multicast (SSM) mode entries.</p> <p>exact—(Optional) Clear only the group that exactly matches the specified group address.</p> <p>inet inet6—(Optional) Clear the PIM entries for IPv4 or IPv6 family addresses, respectively.</p> <p>instance instance-name—(Optional) Clear the entries for a specific PIM-enabled routing instance.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>rp ip-address/prefix source ip-address/prefix—(Optional) Clear the PIM entries with a specified rendezvous point (RP) address and prefix or with a specified source address and prefix. You can omit the prefix.</p> <p>sg star-g—(Optional) Clear PIM (S,G) or (*,G) entries.</p>

Additional Information The `clear pim join` command cannot be used to clear the PIM join and prune state on a backup Routing Engine when nonstop active routing is enabled.

Required Privilege Level clear

Related Documentation • [show pim join on page 512](#)

List of Sample Output [clear pim join on page 386](#)
[clear pim join inet6 on page 386](#)
[clear pim join inet6 star-g on page 386](#)

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

Sample Output

`clear pim join`

```
user@host> clear pim join
Cleared 8 Join/Prune states
```

`clear pim join inet6`

```
user@host> clear pim join inet6
Cleared 4 Join/Prune states
```

`clear pim join inet6 star-g`

```
user@host> clear pim join inet6 star-g
Cleared 1 Join/Prune states
```

clear pim join-distribution

Syntax	clear pim join-distribution <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced in Junos OS Release 10.0.
Description	<p>Redistribute the Protocol Independent Multicast (PIM) join states.</p> <p>You can find out if there are multiple paths available for a source (for example, an RP) with the output of the show pim source command.</p> <p>When you include the join-load-balance statement in the configuration, the PIM join states are distributed evenly on available equal-cost multipath links. When an upstream neighbor link fails, Junos OS redistributes the PIM join states to the remaining links. However, when new links are added or the failed link is restored, the existing PIM joins are not redistributed to the new link. New flows will be distributed to the new links. However, in a network without new joins and prunes, the new link is not used for multicast traffic. The clear pim join-distribution command redistributes the existing flows to the new upstream neighbors. Redistributing the existing flows causes traffic to be disrupted, so we recommend that you run the clear pim join-distribution command during a maintenance window.</p>
Options	<p>none—Redistribute the PIM join states for the default master instance.</p> <p>instance <i>instance-name</i>—(Optional) Redistribute the join states for a specific PIM-enabled routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Additional Information	The clear pim join-distribution command cannot be used to redistribute the PIM join states on a backup Routing Engine when nonstop active routing is enabled.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show pim neighbors on page 541 • show pim join on page 512 • <i>join-load-balance</i>
List of Sample Output	clear pim join-distribution on page 388
Output Fields	When you enter this command, you are provided no feedback on the status of your request. You can enter the show pim join command before and after distributing the join state to verify the operation.

Sample Output

clear pim join-distribution

```
user@host> clear pim join-distribution
```

clear pim register

Syntax	clear pim register <inet inet6> <instance <i>instance-name</i> > <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch and the QFX Series)	clear pim register <inet inet6> <instance <i>instance-name</i> > <interface <i>interface-name</i> >
Syntax (PTX Series)	clear pim register <inet inet6> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced in Junos OS Release 7.6. Command introduced in Junos OS Release 9.0 for EX Series switches. inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Clear Protocol Independent Multicast (PIM) register message counters.
Options	<p>none—Clear PIM register message counters for all family addresses, instances, and interfaces.</p> <p>inet inet6—(Optional) Clear PIM register message counters for IPv4 or IPv6 family addresses, respectively.</p> <p>instance <i>instance-name</i>—(Optional) Clear register message counters for a specific PIM-enabled routing instance.</p> <p>interface <i>interface-name</i>—(Optional) Clear PIM register message counters for a specific interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Additional Information	The clear pim register command cannot be used to clear the PIM register state on a backup Routing Engine when nonstop active routing is enabled.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show pim statistics on page 572
List of Sample Output	clear pim register on page 390
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear pim register

```
user@host> clear pim register
```


clear pim snooping join

Syntax	clear pim snooping join <instance <i>instance-name</i> > <vlan-id <i>vlan-id</i> >
Release Information	Command introduced in Junos OS Release 12.3 for MX Series 3D Universal Edge devices. Command introduced in Junos OS Release 13.2 for M Series Multiservice Edge devices.
Description	Clear information about Protocol Independent Multicast (PIM) snooping joins.
Options	<p>none—Display detailed information.</p> <p>instance <i>instance-name</i>—(Optional) Clear PIM snooping join information for the specified routing instance.</p> <p>vlan-id <i>vlan-identifier</i>—(Optional) Clear PIM snooping join information for the specified VLAN.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>PIM Snooping for VPLS</i>
List of Sample Output	clear pim snooping join on page 391
Output Fields	See show pim snooping join for an explanation of the output fields.

Sample Output

clear pim snooping join

The following sample output displays information about PIM snooping joins before and after the **clear pim snooping join** command is entered:

```

user@host> show pim snooping join extensive
Instance: vpls1
Learning-Domain: vlan-id 10
Learning-Domain: vlan-id 20

Group: 225.1.1.2
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 20.0.120.5, port: ge-1/3/7.20
Downstream port: ge-1/3/1.20
Downstream neighbors:
20.0.120.2 State: Join Flags: SRW Timeout: 185

Group: 225.1.1.3
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 20.0.120.4, port: ge-1/3/5.20

```

```
Downstream port: ge-1/3/3.20
Downstream neighbors:
20.0.120.3 State: Join Flags: SRW Timeout: 175

user@host> clear pim snooping join
Clearing the Join/Prune state for 224.0.0.0/4
Clearing the Join/Prune state for 224.0.0.0/4

user@host> show pim snooping join extensive
Instance: vpls1
Learning-Domain: vlan-id 10
Learning-Domain: vlan-id 20
```

clear pim snooping statistics

Syntax	clear pim snooping statistics <instance <i>instance-name</i> > <interface <i>interface-name</i> > <vlan-id <i>vlan-id</i> >
Release Information	Command introduced in Junos OS Release 12.3 for MX Series 3D Universal Edge devices. Command introduced in Junos OS Release 13.2 for M Series Multiservice Edge devices.
Description	Clear Protocol Independent Multicast (PIM) snooping statistics.
Options	<p>none—Clear PIM snooping statistics for all family addresses, instances, and interfaces.</p> <p>instance <i>instance-name</i>—(Optional) Clear statistics for a specific PIM-snooping-enabled routing instance.</p> <p>interface <i>interface-name</i>—(Optional) Clear PIM snooping statistics for a specific interface.</p> <p>vlan-id <i>vlan-identifier</i>—(Optional) Clear PIM snooping statistics information for the specified VLAN.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • <i>PIM Snooping for VPLS</i>
List of Sample Output	clear pim snooping statistics on page 393
Output Fields	See show pim snooping statistics for an explanation of the output fields.

Sample Output

clear pim snooping statistics

The following sample output displays PIM snooping statistics before and after the **clear pim snooping statistics** command is entered:

```
user@host> show pim snooping statistics
Instance: vpls1
Learning-Domain: vlan-id 10

Tx J/P messages 0
Rx J/P messages 660
Rx J/P messages -- seen 0
Rx J/P messages -- received 660
Rx Hello messages 1396
Rx Version Unknown 0
Rx Neighbor Unknown 0
Rx Upstream Neighbor Unknown 0
Rx Bad Length 0
Rx J/P Busy Drop 0
Rx J/P Group Aggregate 0
Rx Malformed Packet 0
```

```
Learning-Domain: vlan-id 20
user@host> clear pim snooping statistics
user@host> show pim snooping statistics
Instance: vpls1
Learning-Domain: vlan-id 10

Tx J/P messages 0
RX J/P messages 0
Rx J/P messages -- seen 0
Rx J/P messages -- received 0
Rx Hello messages 0
Rx Version Unknown 0
Rx Neighbor Unknown 0
Rx Upstream Neighbor Unknown 0
Rx Bad Length 0
Rx J/P Busy Drop 0
Rx J/P Group Aggregate 0
Rx Malformed Packet 0

Learning-Domain: vlan-id 20
```

clear pim statistics

Syntax	clear pim statistics <inet inet6> <instance <i>instance-name</i> > <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch and the QFX Series)	clear pim statistics <inet inet6> <instance <i>instance-name</i> > <interface <i>interface-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Clear Protocol Independent Multicast (PIM) statistics.
Options	<p>none—Clear PIM statistics for all family addresses, instances, and interfaces.</p> <p>inet inet6—(Optional) Clear PIM statistics for IPv4 or IPv6 family addresses, respectively.</p> <p>instance <i>instance-name</i>—(Optional) Clear statistics for a specific PIM-enabled routing instance.</p> <p>interface <i>interface-name</i>—(Optional) Clear PIM statistics for a specific interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Additional Information	The clear pim statistics command cannot be used to clear the PIM statistics on a backup Routing Engine when nonstop active routing is enabled.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show pim statistics on page 572
List of Sample Output	clear pim statistics on page 395
Output Fields	See show pim statistics for an explanation of output fields.

Sample Output

clear pim statistics

The following sample output displays PIM statistics before and after the **clear pim statistics** command is entered:

```

user@host> show pim statistics
PIM statistics on all interfaces:
PIM Message type      Received      Sent  Rx errors
Hello                  0             0      0
Register               0             0      0
Register Stop         0             0      0
Join Prune             0             0      0
Bootstrap              0             0      0
Assert                0             0      0
Graft                  0             0      0
Graft Ack              0             0      0
Candidate RP           0             0      0
V1 Query               2111          4222     0
V1 Register            0             0      0
V1 Register Stop       0             0      0
V1 Join Prune          14200         13115     0
V1 RP Reachability     0             0      0
V1 Assert              0             0      0
V1 Graft               0             0      0
V1 Graft Ack           0             0      0
PIM statistics summary for all interfaces:
Unknown type           0
V1 Unknown type        0
Unknown Version        0
Neighbor unknown       0
Bad Length             0
Bad Checksum           0
Bad Receive If         0
Rx Intf disabled       2007
Rx V1 Require V2       0
Rx Register not RP     0
RP Filtered Source     0
Unknown Reg Stop       0
Rx Join/Prune no state 1040
Rx Graft/Graft Ack no state 0
...

```

```

user@host> clear pim statistics
user@host> show pim statistics
PIM statistics on all interfaces:
PIM Message type      Received      Sent  Rx errors
Hello                  0             0      0
Register               0             0      0
Register Stop         0             0      0
Join Prune             0             0      0
Bootstrap              0             0      0
Assert                0             0      0
Graft                  0             0      0
Graft Ack              0             0      0
Candidate RP           0             0      0
V1 Query               1             0      0
V1 Register            0             0      0
...

```

request pim multicast-tunnel rebalance

Syntax	request pim multicast-tunnel rebalance <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	request pim multicast-tunnel rebalance <instance <i>instance-name</i> >
Release Information	Command introduced in Junos OS Release 10.2. Command introduced in Junos OS Release 10.2 for EX Series switches.
Description	Rebalance the assignment of multicast tunnel encapsulation interfaces across available tunnel-capable PICs or across a configured list of tunnel-capable PICs. You can determine whether a rebalance is necessary by running the show pim interfaces instance <i>instance-name</i> command.
Options	<p>none—Re-create and rebalance all tunnel interfaces for all routing instances.</p> <p>instance <i>instance-name</i>—Re-create and rebalance all tunnel interfaces for a specific instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	maintenance
Related Documentation	<ul style="list-style-type: none"> • show pim interfaces on page 509 • <i>Load Balancing Multicast Tunnel Interfaces Among Available PICs</i>
Output Fields	This command produces no output. To verify the operation of the command, run the show pim interface instance <i>instance-name</i> before and after running the request pim multicast-tunnel rebalance command.

show amt statistics

Syntax	show amt statistics <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced in JUNOS Release 10.2.
Description	Display information about the Automatic Multicast Tunneling (AMT) protocol tunnel statistics.
Options	<p>none—Display summary information about all AMT Protocol tunnels.</p> <p>instance <i>instance-name</i>—(Optional) Display information for the specified instance only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear amt statistics on page 362 • show amt summary on page 401 • show amt tunnel on page 403
List of Sample Output	show amt statistics on page 399
Output Fields	Table 26 on page 398 describes the output fields for the show amt statistics command. Output fields are listed in the approximate order in which they appear.

Table 26: show amt statistics Output Fields

Field Name	Field Description
AMT receive message count	<p>Summary of AMT statistics for messages received on all interfaces.</p> <ul style="list-style-type: none"> • AMT relay discovery—Number of AMT relay discovery messages received. • AMT membership request—Number of AMT membership request messages received. • AMT membership update—Number of AMT membership update messages received.
AMT send message count	<p>Summary of AMT statistics for messages sent on all interfaces.</p> <ul style="list-style-type: none"> • AMT relay advertisement—Number of AMT relay advertisement messages sent. • AMT membership query—Number of AMT membership query messages sent.

Table 26: show amt statistics Output Fields (*continued*)

Field Name	Field Description
AMT error message count	<p>Summary of AMT statistics for error messages received on all interfaces.</p> <ul style="list-style-type: none"> • AMT incomplete packet—Number of messages received with length errors so severe that further classification could not occur. • AMT invalid mac—Number of messages received with an invalid message authentication code (MAC). • AMT unexpected type—Number of messages received with an unknown message type specified. • AMT invalid relay discovery address—Number of AMT relay discovery messages received with an address other than the configured anycast address. • AMT invalid membership request address—Number of AMT membership request messages received with an address other than the configured AMT local address. • AMT invalid membership update address—Number of AMT membership update messages received with an address other than the configured AMT local address. • AMT incomplete relay discovery messages—Number of AMT relay discovery messages received that are not fully formed. • AMT incomplete membership request messages—Number of AMT membership request messages received that are not fully formed. • AMT incomplete membership update messages—Number of AMT membership update messages received that are not fully formed. • AMT no active gateway—Number of AMT membership update messages received for a tunnel that does not exist for the gateway that sent the message. • AMT invalid inner header checksum—Number of AMT membership update messages received with an invalid IP checksum. • AMT gateways timed out—Number of gateways that timed out because of inactivity.

Sample Output

show amt statistics

```
user@host> show amt statistics
```

```

AMT receive message count
AMT relay advertisement           :           2
AMT membership request           :           5
AMT membership update            :           5

AMT send message count
AMT relay advertisement           :           2
AMT membership query             :           5

AMT error message count
AMT incomplete packet             :           0
AMT invalid mac                   :           0
AMT unexpected type               :           0
AMT invalid relay discovery address :           0
AMT invalid membership request address :           0
AMT invalid membership update address :           0
AMT incomplete relay discovery messages :           0
AMT incomplete membership request messages :           0
AMT incomplete membership update messages :           0
AMT no active gateway             :           0

```

AMT invalid inner header checksum	:	0
AMT gateways timed out	:	0

show amt summary

Syntax	show amt summary <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced in JUNOS Release 10.2.
Description	Display summary information about the Automatic Multicast Tunneling (AMT) protocol.
Options	<p>none—Display summary information about all AMT protocol instances.</p> <p>instance <i>instance-name</i>—(Optional) Display information for the specified instance only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear amt tunnel on page 363 • show amt statistics on page 398 • show amt tunnel on page 403
List of Sample Output	show amt summary on page 402
Output Fields	Table 27 on page 401 describes the output fields for the show amt summary command. Output fields are listed in the approximate order in which they appear.

Table 27: show amt summary Output Fields

Field Name	Field Description	Level of Output
AMT anycast prefix	Prefix advertised by unicast routing protocols to route AMT discovery messages to the router from nearby AMT gateways.	All levels
AMT anycast address	Anycast address configured from which the anycast prefix is derived.	All levels
AMT local address	Local unique AMT relay IP address configured. Used to send AMT relay advertisement messages, it is the IP source address of AMT control messages and the source address of the data tunnel encapsulation.	All levels
AMT tunnel limit	Maximum number of AMT tunnels that can be created.	All levels
active tunnels	Number of active AMT tunnel interfaces.	All levels

Sample Output

show amt summary

```
user@host> show amt summary
  AMT anycast prefix : 20.0.0.4/32
  AMT anycast address : 20.0.0.4
  AMT local address : 20.0.0.4
  AMT tunnel limit : 1000, active tunnels : 2
```

show amt tunnel

Syntax	<pre>show amt tunnel <brief detail> <gateway-address <i>gateway-ip-address</i>> <port <i>port-number</i>> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <tunnel-interface <i>interface-name</i>></pre>	
Release Information	Command introduced in JUNOS Release 10.2.	
Description	Display information about the Automatic Multicast Tunneling (AMT) dynamic tunnels.	
Options	<p>none—Display summary information about all AMT protocol instances.</p> <p>brief detail—(Optional) Display the specified level of detail.</p> <p>gateway-address <i>gateway-ip-address</i> port <i>port-number</i>—(Optional) Display information for the specified AMT gateway only. If no port is specified, display information for all AMT gateways with the given IP address.</p> <p>instance <i>instance-name</i>—(Optional) Display information for the specified instance only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>tunnel-interface <i>interface-name</i>—(Optional) Display information for the specified AMT tunnel interface only.</p>	
Required Privilege Level	view	
Related Documentation	<ul style="list-style-type: none"> • clear amt tunnel on page 363 • show amt statistics on page 398 • show amt summary on page 401 	
List of Sample Output	show amt tunnel on page 404 show amt tunnel detail on page 404 show amt tunnel tunnel-interface on page 405 show amt tunnel gateway-address on page 405 show amt tunnel gateway-address detail on page 405	
Output Fields	<p>Table 28 on page 403 describes the output fields for the show amt tunnel command. Output fields are listed in the approximate order in which they appear.</p>	

Table 28: show amt tunnel Output Fields

Field Name	Field Description	Level of Output
AMT gateway address	Address of the AMT gateway that is being connected by the AMT tunnel.	All levels

Table 28: show amt tunnel Output Fields (*continued*)

Field Name	Field Description	Level of Output
port	Client port used by the AMT tunnel.	All levels
AMT tunnel interface	Dynamically created AMT logical interfaces used by the AMT tunnel in the format ud-FPC/PIC/Port.unit .	All levels
AMT tunnel state	State of the AMT tunnel. The state is normally Active . <ul style="list-style-type: none"> • Active—The tunnel is active. • Pending—The tunnel creation is pending. This is a transient state. • Down—The tunnel is in the down state. • Graceful restart pending—Graceful restart is in progress. • Reviving—The routing protocol daemon or Routing Engine was restarted (not gracefully). The tunnel remains in the reviving state until the AMT gateway sends a control message. When the message is received the tunnel is moved to the Active state. If no message is received before the AMT tunnel inactivity timer expires, the tunnel is deleted. 	All levels
AMT tunnel inactivity timeout	Number of seconds since the most recent control message was received from an AMT gateway. If no message is received before the AMT tunnel inactivity timer expires, the tunnel is deleted.	All levels
Number of groups	Number of multicast groups using the tunnel.	All levels
Group	Multicast group address or addresses using the tunnel.	detail
Include Source	Multicast source address for each IGMPv3 group using the tunnel.	detail
AMT message count	Statistics for AMT messages: <ul style="list-style-type: none"> • AMT Request—Number of AMT relay tunnel request messages received. • AMT membership update—Number of AMT membership update messages received. 	All levels

Sample Output

show amt tunnel

```

user@host> show amt tunnel
AMT gateway address : 11.11.11.2, port : 2268
AMT tunnel interface : ud-5/1/10.1120256
AMT tunnel state : Active
AMT tunnel inactivity timeout : 15
Number of groups : 1

AMT message count:
AMT Request      AMT membership update
2                2

```

show amt tunnel detail

```

user@host> show amt tunnel detail
AMT gateway address : 11.11.11.2, port : 2268
AMT tunnel interface : ud-5/3/10.1120512

```

```

AMT tunnel state : Active
AMT tunnel inactivity timeout : 62
Number of groups : 1
Group: 226.2.3.2

AMT message count:
AMT Request      AMT membership update
2                2

AMT gateway address : 11.11.11.3, port : 2268
AMT tunnel interface : ud-5/2/10.1120513
AMT tunnel state : Active
AMT tunnel inactivity timeout : 214
Number of groups : 1
Group: 226.2.3.3

AMT message count:
AMT Request      AMT membership update
2                2

```

show amt tunnel tunnel-interface

```

user@host> show amt tunnel tunnel-interface ud-5/3/10.1120512
AMT gateway address : 11.11.11.2, port : 2268
AMT tunnel interface : ud-5/3/10.1120512
AMT tunnel state : Active
AMT tunnel inactivity timeout : 145
Number of groups : 1

AMT message count:
AMT Request      AMT membership update
2                2

```

show amt tunnel gateway-address

```

user@host> show amt tunnel gateway-address 11.11.11.3 port 2268
AMT gateway address : 11.11.11.3, port : 2268
AMT tunnel interface : ud-5/2/10.1120513
AMT tunnel state : Active
AMT tunnel inactivity timeout : 214
Number of groups : 1
Group: 226.2.3.3

AMT message count:
AMT Request      AMT membership update
2                2

```

show amt tunnel gateway-address detail

```

user@host> show amt tunnel gateway-address 11.11.11.2 detail
AMT gateway address : 11.11.11.2, port : 2268
AMT tunnel interface : ud-5/3/10.1120512
AMT tunnel state : Active
AMT tunnel inactivity timeout : 234
Number of groups : 1
Group: 226.2.3.2

AMT message count:
AMT Request      AMT membership update
2                2

```

show dvmrp interfaces

Syntax	show dvmrp interfaces <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display information about Distance Vector Multicast Routing Protocol (DVMRP)–enabled interfaces.
Options	<p>none—(Same as logical-system all) Display information about DVMRP-enabled interfaces.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show dvmrp interfaces on page 407
Output Fields	Table 29 on page 406 describes the output fields for the show dvmrp interfaces command. Output fields are listed in the approximate order in which they appear.

Table 29: show dvmrp interfaces Output Fields

Field Name	Field Description
Interface	Name of the interface.
State	State of the interface: up or down .
Leaf	Whether the interface is a leaf (that is, whether it has no neighbors) or whether it has neighbors.
Metric	Interface metric: a value from 1 through 31.
Announce	Number of routes the interface is announcing.
Mode	DVMRP mode: <ul style="list-style-type: none"> • Forwarding—DVMRP does both the routing and the multicast data forwarding. • Unicast-routing—DVMRP does only the routing. Forwarding of the multicast data packets can be done by enabling PIM on the interface.

Sample Output

show dvmrp interfaces

```
user@host> show dvmrp interfaces
Interface State Leaf Metric Announce Mode
fxp0.0    Up    N    1    4 Forwarding
fxp1.0    Up    N    1    4 Forwarding
fxp2.0    Up    N    1    3 Forwarding
lo0.0     Up    Y    1    0 Unicast-routing
```

show dvmrp neighbors

Syntax	show dvmrp neighbors <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display information about Distance Vector Multicast Routing Protocol (DVMRP) neighbors.
Options	<p>none—(Same as logical-system all) Display information about DVMRP neighbors.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show dvmrp neighbors on page 409
Output Fields	<p>Table 30 on page 408 describes the output fields for the show dvmrp neighbors command. Output fields are listed in the approximate order in which they appear.</p>

Table 30: show dvmrp neighbors Output Fields

Field Name	Field Description
Neighbor	Address of the neighboring DVMRP router.
Interface	Interface through which the neighbor is reachable.
Version	Version of DVMRP that the neighbor is running, in the format <i>majorminor</i> .
Flags	<p>Information about the neighbor:</p> <ul style="list-style-type: none"> 1—One way. The local router has seen the neighbor, but the neighbor has not seen the local router. G—Neighbor supports generation ID. L—Neighbor is a leaf router. M—Neighbor supports mtrace. N—Neighbor supports netmask in prune messages and graft messages. P—Neighbor supports pruning. S—Neighbor supports SNMP.
Routes	Number of routes learned from the neighbor.
Timeout	How long until the DVMRP neighbor information times out, in seconds.
Transitions	Number of generation ID changes that have occurred since the local router learned about the neighbor.

Sample Output

show dvmrp neighbors

```
user@host> show dvmrp neighbors
Neighbor      Interface      Version  Flags    Routes  Timeout  Transitions
192.168.1.1    ipip.0         3.255    PGM      3       28       1
```

show dvmrp prefix

Syntax	show dvmrp prefix <brief detail> <logical-system (all <i>logical-system-name</i>)> <prefix>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display information about Distance Vector Multicast Routing Protocol (DVMRP) prefixes.
Options	<p>none—Display standard information about all DVMRP prefixes.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>prefix—(Optional) Display information about specific prefixes.</p>
Required Privilege Level	view
List of Sample Output	show dvmrp prefix on page 411 show dvmrp prefix brief on page 411 show dvmrp prefix detail on page 411
Output Fields	Table 31 on page 410 describes the output fields for the show dvmrp prefix command. Output fields are listed in the approximate order in which they appear.

Table 31: show dvmrp prefix Output Fields

Field Name	Field Description	Level of Output
Prefix	DVMRP route.	All levels
Next hop	Next hop from which the route was learned.	All levels
Age	Last time that the route was refreshed.	All levels
<i>multicast-group</i>	Multicast group address.	detail
Prunes sent	Number of prune messages sent to the multicast group.	detail
Grafts sent	Number of grafts sent to the multicast group.	detail
Cache lifetime	Lifetime of the group in the multicast cache, in seconds.	detail
Prune lifetime	Lifetime remaining and total lifetime of prune messages, in seconds.	detail

Sample Output

show dvmrp prefix

```
user@host> show dvmrp prefix
Prefix          Next hop      Age
10.38.0.0       /30 10.38.0.1 00:06:17
10.38.0.4       /30 10.38.0.5 00:06:13
10.38.0.8       /30 10.38.0.2 00:00:04
10.38.0.12      /30 10.38.0.6 00:00:04
10.255.14.114   /32 10.255.14.114 00:06:17
10.255.14.142   /32 10.38.0.2 00:00:04
10.255.14.144   /32 10.38.0.2 00:00:04
10.255.70.15    /32 10.38.0.6 00:00:04
192.168.14.0    /24 192.168.14.114 00:06:17
192.168.195.40  /30 192.168.195.41 00:06:17
192.168.195.92  /30 10.38.0.2 00:00:04
```

show dvmrp prefix brief

The output for the **show dvmrp prefix brief** command is identical to that for the **show dvmrp prefix** command.

show dvmrp prefix detail

```
user@host> show dvmrp prefix detail
Prefix          Next hop      Age
10.38.0.0       /30 10.38.0.1 00:06:28
10.38.0.4       /30 10.38.0.5 00:06:24
10.38.0.8       /30 10.38.0.2 00:00:15
10.38.0.12      /30 10.38.0.6 00:00:15
10.255.14.114   /32 10.255.14.114 00:06:28
10.255.14.142   /32 10.38.0.2 00:00:15
10.255.14.144   /32 10.38.0.2 00:00:15
10.255.70.15    /32 10.38.0.6 00:00:15
192.168.14.0    /24 192.168.14.114 00:06:28
192.168.195.40  /30 192.168.195.41 00:06:28
192.168.195.92  /30 10.38.0.2 00:00:15
```

show dvmrp prunes

Syntax	show dvmrp prunes <all rx tx> <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display information about active Distance Vector Multicast Routing Protocol (DVMRP) prune messages.
Options	<p>none—Display received and transmitted DVMRP prune information.</p> <p>all—(Optional) Display information about all received and transmitted prune messages.</p> <p>rx—(Optional) Display information about received prune messages.</p> <p>tx—(Optional) Display information about transmitted prune messages.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show dvmrp prunes on page 412
Output Fields	Table 32 on page 412 describes the output fields for the show dvmrp prunes command. Output fields are listed in the approximate order in which they appear.

Table 32: show dvmrp prunes Output Fields

Field Name	Field Description
Group	Group address.
Source prefix	Prefix for the prune.
Timeout	How long until the prune message expires, in seconds.
Neighbor	Neighbor to which the prune was sent or from which the prune was received.

Sample Output

show dvmrp prunes

```

user@host> show dvmrp prunes
Group          Source prefix      Timeout Neighbor
224.0.1.1      128.112.0.0       /12    7077 192.168.1.1
224.0.1.32     160.0.0.0         /3     7087 192.168.1.1
224.2.123.4    136.0.0.0         /5     6955 192.168.1.1
224.2.127.1    129.0.0.0         /8     7046 192.168.1.1

```

```
224.2.135.86 128.102.128.0 /17 7071 192.168.1.1
224.2.135.86 129.0.0.0 /8 7074 192.168.1.1
224.2.135.86 130.0.0.0 /7 7071 192.168.1.1
...
```

show igmp group

Syntax	show igmp group <brief detail> <group-name> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch and the QFX Series)	show igmp group <brief detail> <group-name>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display Internet Group Management Protocol (IGMP) group membership information.
Options	<p>none—Display standard information about membership for all IGMP groups.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>group-name—(Optional) Display group membership for the specified IP address only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> clear igmp membership on page 364
List of Sample Output	show igmp group (Include Mode) on page 415 show igmp group (Exclude Mode) on page 416 show igmp group brief on page 416 show igmp group detail on page 416
Output Fields	<p>Table 33 on page 414 describes the output fields for the show igmp group command. Output fields are listed in the approximate order in which they appear.</p>

Table 33: show igmp group Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface that received the IGMP membership report. A name of local indicates that the local routing device joined the group itself.	All levels
Group	Group address.	All levels
Group Mode	Mode the SSM group is operating in: Include or Exclude .	All levels
Source	Source address.	All levels

Table 33: show igmp group Output Fields (*continued*)

Field Name	Field Description	Level of Output
Source timeout	Time remaining until the group traffic is no longer forwarded. The timer is refreshed when a listener in include mode sends a report. A group in exclude mode or configured as a static group displays a zero timer.	detail
Last reported by	Address of the host that last reported membership in this group.	All levels
Timeout	Time remaining until the group membership is removed.	brief none
Group timeout	Time remaining until a group in exclude mode moves to include mode. The timer is refreshed when a listener in exclude mode sends a report. A group in include mode or configured as a static group displays a zero timer.	detail
Type	Type of group membership: <ul style="list-style-type: none"> • Dynamic—Host reported the membership. • Static—Membership is configured. 	All levels

Sample Output

show igmp group (Include Mode)

```

user@host> show igmp group
Interface: t1-0/1/0.0
  Group: 232.1.1.1
    Group mode: Include
    Source: 10.0.0.2
    Last reported by: 10.9.5.2
    Timeout:      24 Type: Dynamic
  Group: 232.1.1.1
    Group mode: Include
    Source: 10.0.0.3
    Last reported by: 10.9.5.2
    Timeout:      24 Type: Dynamic
  Group: 232.1.1.1
    Group mode: Include
    Source: 10.0.0.4
    Last reported by: 10.9.5.2
    Timeout:      24 Type: Dynamic
  Group: 232.1.1.2
    Group mode: Include
    Source: 10.0.0.4
    Last reported by: 10.9.5.2
    Timeout:      24 Type: Dynamic
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
  Group: 224.0.0.2
    Source: 0.0.0.0
    Last reported by: Local
    Timeout:      0 Type: Dynamic
  Group: 224.0.0.22
    Source: 0.0.0.0

```

```
Last reported by: Local
Timeout:          0 Type: Dynamic
```

show igmp group (Exclude Mode)

```
user@host> show igmp group
Interface: t1-0/1/0.0
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
  Group: 224.0.0.2
    Source: 0.0.0.0
    Last reported by: Local
    Timeout:          0 Type: Dynamic
  Group: 224.0.0.22
    Source: 0.0.0.0
    Last reported by: Local
    Timeout:          0 Type: Dynamic
```

show igmp group brief

The output for the **show igmp group brief** command is identical to that for the **show igmp group** command.

show igmp group detail

```
user@host> show igmp group detail
Interface: t1-0/1/0.0
  Group: 232.1.1.1
    Group mode: Include
    Source: 10.0.0.2
    Source timeout: 12
    Last reported by: 10.9.5.2
    Group timeout:          0 Type: Dynamic
  Group: 232.1.1.1
    Group mode: Include
    Source: 10.0.0.3
    Source timeout: 12
    Last reported by: 10.9.5.2
    Group timeout:          0 Type: Dynamic
  Group: 232.1.1.1
    Group mode: Include
    Source: 10.0.0.4
    Source timeout: 12
    Last reported by: 10.9.5.2
    Group timeout:          0 Type: Dynamic
  Group: 232.1.1.2
    Group mode: Include
    Source: 10.0.0.4
    Source timeout: 12
    Last reported by: 10.9.5.2
    Group timeout:          0 Type: Dynamic
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
  Group: 224.0.0.2
    Group mode: Exclude
    Source: 0.0.0.0
    Source timeout: 0
```

```
      Last reported by: Local
      Group timeout:      0 Type: Dynamic
Group: 224.0.0.22
      Group mode: Exclude
      Source: 0.0.0.0
      Source timeout: 0
      Last reported by: Local
      Group timeout:      0 Type: Dynamic
```

show igmp interface

Syntax	show igmp interface <brief detail> <interface-name> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and the QFX Series)	show igmp interface <brief detail> <interface-name>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display information about Internet Group Management Protocol (IGMP)-enabled interfaces.
Options	<p>none—Display standard information about all IGMP-enabled interfaces.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>interface-name—(Optional) Display information about the specified IGMP-enabled interface only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> clear igmp membership on page 364
List of Sample Output	show igmp interface on page 420 show igmp interface brief on page 420 show igmp interface detail on page 421 show igmp interface <interface-name> on page 421
Output Fields	Table 34 on page 418 describes the output fields for the show igmp interface command. Output fields are listed in the approximate order in which they appear.

Table 34: show igmp interface Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface.	All levels
Querier	Address of the routing device that has been elected to send membership queries.	All levels
State	State of the interface: Up or Down .	All levels

Table 34: show igmp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
SSM Map Policy	Name of the source-specific multicast (SSM) map policy that has been applied to the IGMP interface.	All levels
Timeout	How long until the IGMP querier is declared to be unreachable, in seconds.	All levels
Version	IGMP version being used on the interface: 1, 2, or 3.	All levels
Groups	Number of groups on the interface.	All levels
Group limit	Maximum number of groups allowed on the interface. Any joins requested after the limit is reached are rejected.	All levels
Group threshold	Configured threshold at which a warning message is generated. This threshold is based on a percentage of groups received on the interface. If the number of groups received reaches the configured threshold, the device generates a warning message.	All levels
Group log-interval	Time (in seconds) between consecutive log messages.	All levels
Immediate Leave	State of the immediate leave option: <ul style="list-style-type: none"> • On—Indicates that the router removes a host from the multicast group as soon as the router receives a leave group message from a host associated with the interface. • Off—Indicates that after receiving a leave group message, instead of removing a host from the multicast group immediately, the router sends a group query to determine if another receiver responds. 	All levels
Promiscuous Mode	State of the promiscuous mode option: <ul style="list-style-type: none"> • On—Indicates that the router can accept IGMP reports from subnetworks that are not associated with its interfaces. • Off—Indicates that the router can accept IGMP reports only from subnetworks that are associated with its interfaces. 	All levels
Passive	State of the passive mode option: <ul style="list-style-type: none"> • On—Indicates that the router can run IGMP on the interface but not send or receive control traffic such as IGMP reports, queries, and leaves. • Off—Indicates that the router can run IGMP on the interface and send or receive control traffic such as IGMP reports, queries, and leaves. <p>The passive statement enables you to selectively activate up to two out of a possible three available query or control traffic options. When enabled, the following options appear after the on state declaration:</p> <ul style="list-style-type: none"> • send-general-query—The interface sends general queries. • send-group-query—The interface sends group-specific and group-source-specific queries. • allow-receive—The interface receives control traffic. 	All levels
OIF map	Name of the OIF map (if configured) associated with the interface.	All levels

Table 34: show igmp interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
SSM map	Name of the source-specific multicast (SSM) map (if configured) used on the interface.	All levels
Configured Parameters	Information configured by the user: <ul style="list-style-type: none"> • IGMP Query Interval—Interval (in seconds) at which this router sends membership queries when it is the querier. • IGMP Query Response Interval—Time (in seconds) that the router waits for a report in response to a general query. • IGMP Last Member Query Interval—Time (in seconds) that the router waits for a report in response to a group-specific query. • IGMP Robustness Count—Number of times the router retries a query. 	All levels
Derived Parameters	Derived information: <ul style="list-style-type: none"> • IGMP Membership Timeout—Timeout period (in seconds) for group membership. If no report is received for these groups before the timeout expires, the group membership is removed. • IGMP Other Querier Present Timeout—Time (in seconds) that the router waits for the IGMP querier to send a query. 	All levels

Sample Output

show igmp interface

```

user@host> show igmp interface
Interface: at-0/3/1.0
  Querier: 10.111.30.1
  State:      Up Timeout:   None Version:  2 Groups:    4
  SSM Map Policy: ssm-policy-A
Interface: so-1/0/0.0
  Querier: 10.111.10.1
  State:      Up Timeout:   None Version:  2 Groups:    2
  SSM Map Policy: ssm-policy-B
Interface: so-1/0/1.0
  Querier: 10.111.20.1
  State:      Up Timeout:   None Version:  2 Groups:    4
  SSM Map Policy: ssm-policy-C
Immediate Leave: On
Promiscuous Mode: Off

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Derived Parameters:
IGMP Membership Timeout: 260.0
IGMP Other Querier Present Timeout: 255.0

```

show igmp interface brief

The output for the **show igmp interface brief** command is identical to that for the **show igmp interface** command. For sample output, see [show igmp interface on page 420](#).

show igmp interface detail

The output for the **show igmp interface detail** command is identical to that for the **show igmp interface** command. For sample output, see [show igmp interface on page 420](#).

show igmp interface <interface-name>

```
user@host# show igmp interface ge-3/2/0.0
Interface: ge-3/2/0.0
Querier: 20.1.1.1
State: Up Timeout:   None Version:  3 Groups:    1
Group limit: 8
Group threshold: 60
Group log-interval: 10
Immediate leave: Off
Promiscuous mode: Off
```

show igmp snooping interface

Syntax	show igmp snooping interface <i>interface-name</i> <brief detail> <bridge-domain <i>bridge-domain-name</i> > <virtual-switch <i>virtual-switch-name</i> > <vlan-id <i>vlan-identifier</i> >
Release Information	Command introduced in Junos OS Release 8.5.
Description	Display IGMP snooping interface information.
Options	<p>none—Display detailed information.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>bridge-domain <i>bridge-domain-name</i>—(Optional) Display information about a particular bridge domain.</p> <p>virtual-switch <i>virtual-switch-name</i>—(Optional) Display information about a particular virtual switch.</p> <p>vlan-id <i>vlan-identifier</i>—(Optional) Display information about a particular VLAN.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show igmp snooping membership on page 425 • show igmp snooping statistics on page 429
List of Sample Output	show igmp snooping interface on page 423 show igmp snooping interface (Group Limit Configured) on page 424
Output Fields	Table 35 on page 422 lists the output fields for the show igmp snooping interface command. Output fields are listed in the approximate order in which they appear.

Table 35: show igmp snooping interface Output Fields

Field Name	Field Description	Level of Output
Routing-instance	Routing instance for IGMP snooping.	All levels
Learning Domain	Learning domain for snooping.	All levels
IGMP Query Interval	Frequency (in seconds) with which this router sends membership queries when it is the querier.	detail
IGMP Query Response Interval	Time (in seconds) that the router waits for a response to a general query.	detail

Table 35: show igmp snooping interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
IGMP Last Member Query Interval	Time (in seconds) that the router waits for a report in response to a group-specific query.	detail
IGMP Robustness Count	Number of times the router retries a query.	detail
immediate-leave	State of immediate leave: On or Off .	All levels
router-interface	Router interfaces that are part of this learning domain.	All levels
Group limit	Maximum number of (source,group) pairs allowed per interface. When a group limit is not configured, this field is not shown.	All levels
interface	Interfaces that are being snooped in this learning domain.	All levels
Groups	Number of groups on the interface.	none
State	State of the interface: Up or Down .	none
Up Groups	Number of active multicast groups attached to the logical interface.	All levels
IGMP Membership Timeout	Timeout for group membership. If no report is received for these groups before the timeout expires, the group membership is removed.	none
IGMP Other Querier Present Timeout	Time that the router waits for the IGMP querier to send a query.	none

Sample Output

show igmp snooping interface

```

user@host> show igmp snooping interface
Instance: bridge-domain bar

Learning-Domain: default
Interface: ge-0/1/0.200
  State:          Up Groups:      0
  Immediate leave: Off
  Router interface: yes
Interface: ge-0/1/2.200
  State:          Up Groups:      2
  Immediate leave: On
  Router interface: no
Interface: ge-0/1/3.200
  State:          Up Groups:      1
  Immediate leave: Off
  Router interface: no

Configured Parameters:
IGMP Query Interval: 130.0
IGMP Query Response Interval: 15.0

```

```
IGMP Last Member Query Interval: 2.0
IGMP Robustness Count: 3

Derived Parameters:
IGMP Membership Timeout: 405.0
IGMP Other Querier Present Timeout: 397.500
```

Sample Output

show igmp snooping interface (Group Limit Configured)

```
user@host> show igmp snooping interface instance vpls1
Instance: vpls1

Learning-Domain: default
Interface: ge-1/3/9.0
  State:          Up Groups:      0
  Immediate leave: Off
  Router interface: yes
Interface: ge-1/3/8.0
  State:          Up Groups:      0
  Immediate leave: Off
  Router interface: yes
  Group limit:    1000

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2
```

show igmp snooping membership

Syntax	show igmp snooping membership <brief detail> <bridge-domain <i>bridge-domain-name</i> > <group <i>group-name</i> > <virtual-switch <i>virtual-switch-name</i> > <vlan-id <i>vlan-identifier</i> >
Release Information	Command introduced in Junos OS Release 8.5.
Description	Display IGMP snooping membership information.
Options	<p>none—Display detailed information.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>bridge-domain <i>bridge-domain-name</i>—(Optional) Display information about a particular bridge domain.</p> <p>group <i>group-name</i> —(Optional) Display information about this group address.</p> <p>virtual-switch <i>virtual-switch-name</i>—(Optional) Display information about a particular virtual switch.</p> <p>vlan-id <i>vlan-identifier</i>—(Optional) Display information about a particular VLAN.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show igmp snooping interface on page 422 • show igmp snooping statistics on page 429 • clear igmp snooping membership on page 367
List of Sample Output	show igmp snooping membership on page 426 show igmp snooping membership (Exclude Mode) on page 427 show igmp snooping membership interface ge-0/1/2.200 on page 427 show igmp snooping membership vlan-id 1 on page 427
Output Fields	Table 36 on page 425 lists the output fields for the show igmp snooping membership command. Output fields are listed in the approximate order in which they appear.

Table 36: show igmp snooping membership Output Fields

Field Name	Field Description	Level of Output
Instance	Routing instance for IGMP snooping.	All levels
Learning Domain	Learning domain for snooping.	All levels

Table 36: show igmp snooping membership Output Fields (*continued*)

Field Name	Field Description	Level of Output
Interface	Interface on which this router is a proxy.	detail
Up Groups	Number of active multicast groups attached to the logical interface.	All levels
Group	Multicast group address in the membership database.	All levels
Group Mode	Mode the SSM group is operating in: Include or Exclude .	All levels
Source	Source address used on queries.	detail
Last reported by	Address of source last replying to the query.	detail
Group Timeout	Time remaining until a group in exclude mode moves to include mode. The timer is refreshed when a listener in exclude mode sends a report. A group in include mode or configured as a static group displays a zero timer.	All levels
Timeout	Length of time (in seconds) left until the entry is purged.	detail
Type	Way that the group membership information was learned: <ul style="list-style-type: none"> • Dynamic—Group membership was learned by the IGMP protocol. • Static—Group membership was learned by configuration. 	detail
Include receiver	Source address of receiver included in membership with timeout (in seconds).	detail

Sample Output

show igmp snooping membership

```

user@host> show igmp snooping membership
Instance: vpls2

Learning-Domain: vlan-id 2
Interface: ge-3/0/0.2
Up Groups:      0
Interface: ge-3/1/0.2
Up Groups:      0
Interface: ge-3/1/5.2
Up Groups:      0

Instance: vpls1

Learning-Domain: vlan-id 1
Interface: ge-3/0/0.1
Up Groups:      0
Interface: ge-3/1/0.1
Up Groups:      0
Interface: ge-3/1/5.1
Up Groups:      1
  Group: 225.10.10.1
    Group mode: Exclude
    Source: 0.0.0.0

```

```

Last reported by: 100.6.85.2
Group timeout:    173 Type: Dynamic

```

show igmp snooping membership (Exclude Mode)

```

user@host> show igmp snooping membership
Instance: vpls2

Learning-Domain: vlan-id 2
Interface: ge-3/0/0.2
Up Groups:      0
Interface: ge-3/1/0.2
Up Groups:      0
Interface: ge-3/1/5.2
Up Groups:      0

Instance: vpls1

Learning-Domain: vlan-id 1
Interface: ge-3/0/0.1
Up Groups:      0
Interface: ge-3/1/0.1
Up Groups:      0
Interface: ge-3/1/5.1
Up Groups:      1
  Group: 225.10.10.1
    Group mode: Exclude
    Source: 0.0.0.0
    Last reported by: 100.6.85.2
    Group timeout:    173 Type: Dynamic

```

show igmp snooping membership interface ge-0/1/2.200

```

user@host> show igmp snooping membership interface ge-0/1/2.200
Instance: bridge-domain bar

Learning-Domain: default
Interface: ge-0/1/2.200
  Group: 225.1.1.1
    Source: 0.0.0.0
    Timeout: 391 Type: Static
  Group: 232.1.1.1
    Source: 192.168.1.1
    Timeout: 0 Type: Static

```

show igmp snooping membership vlan-id 1

```

user@host> show igmp snooping membership vlan-id 1
Instance: vpls2

Instance: vpls1

Learning-Domain: vlan-id 1
Interface: ge-3/0/0.1
Up Groups:      0
Interface: ge-3/1/0.1
Up Groups:      0
Interface: ge-3/1/5.1
Up Groups:      1
  Group: 225.10.10.1
    Group mode: Exclude
    Source: 0.0.0.0

```

Last reported by: 100.6.85.2
Group timeout: 209 Type: Dynamic

show igmp snooping statistics

Syntax	show igmp snooping statistics <brief detail> <bridge-domain <i>bridge-domain-name</i> > <virtual-switch <i>virtual-switch-name</i> > <vlan-id <i>vlan-identifier</i> >
Release Information	Command introduced in Junos OS Release 8.5.
Description	Display IGMP snooping statistics.
Options	<p>none—(Optional) Display detailed information.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>bridge-domain <i>bridge-domain-name</i>—(Optional) Display information about a particular bridge domain.</p> <p>virtual-switch <i>virtual-switch-name</i>—(Optional) Display information about a particular virtual switch.</p> <p>vlan-id <i>vlan-identifier</i>—(Optional) Display information about a particular VLAN.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show igmp snooping interface on page 422 • show igmp snooping membership on page 425 • clear igmp snooping statistics on page 368
List of Sample Output	show igmp snooping statistics on page 430
Output Fields	Table 37 on page 429 lists the output fields for the show igmp snooping statistics command. Output fields are listed in the approximate order in which they appear.

Table 37: show igmp snooping statistics Output Fields

Field Name	Field Description	Level of Output
Routing-instance	Routing instance for IGMP snooping.	All levels
IGMP packet statistics	Heading for IGMP snooping statistics for all interfaces or for the specified interface.	All levels
learning-domain	Appears at end of “IGMP packets statistics” line.	All levels

Table 37: show igmp snooping statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
IGMP Message type	Summary of IGMP statistics: <ul style="list-style-type: none"> • Membership Query—Number of membership queries sent and received. • V1 Membership Report—Number of version 1 membership reports sent and received. • DVMRP—Number of DVMRP messages sent or received. • PIM V1—Number of PIM version 1 messages sent or received. • Cisco Trace—Number of Cisco trace messages sent or received. • V2 Membership Report—Number of version 2 membership reports sent or received. • Group Leave—Number of group leave messages sent or received. • Domain Wide Report—Number of domain-wide reports sent or received. • V3 Membership Report—Number of version 3 membership reports sent or received. • Other Unknown types—Number of unknown message types received. • IGMP v3 unsupported type—Number of messages received with unknown and unsupported IGMP version 3 message types. • IGMP v3 source required for SSM—Number of IGMP version 3 messages received that contained no source. • IGMP v3 mode not applicable for SSM—Number of IGMP version 3 messages received that did not contain a mode applicable for source-specific multicast (SSM). 	All levels
Received	Number of messages received.	All levels
Sent	Number of messages sent.	All levels
Rx errors	Number of received packets that contained errors.	All levels
IGMP Global Statistics	Summary of IGMP snooping statistics for all interfaces. <ul style="list-style-type: none"> • Bad Length—Number of messages received with length errors so severe that further classification could not occur. • Bad Checksum—Number of messages received with a bad IP checksum. No further classification was performed. • Rx non-local—Number of messages received from senders that are not local. 	All levels

Sample Output

show igmp snooping statistics

```
user@host> show igmp snooping statistics
Routing-instance foo
```

```
IGMP packet statistics for all interfaces in learning-domain vlan-100
```

IGMP Message type	Received	Sent	Rx errors
Membership Query	89	51	0
V1 Membership Report	0	0	0
DVMRP	0	0	0

PIM V1	0	0	0
Cisco Trace	0	0	0
V2 Membership Report	139	0	0
Group Leave	0	0	0
Domain Wide Report	0	0	0
V3 Membership Report	136	0	0
Other Unknown types			0
IGMP v3 unsupported type			0
IGMP v3 source required for SSM			23
IGMP v3 mode not applicable for SSM			0

IGMP Global Statistics

Bad Length	0
Bad Checksum	0
Rx non-local	0

Routing-instance bar

IGMP packet statistics for all interfaces in learning-domain vlan-100

IGMP Message type	Received	Sent	Rx errors
Membership Query	89	51	0
V1 Membership Report	0	0	0
DVMRP	0	0	0
PIM V1	0	0	0
Cisco Trace	0	0	0
V2 Membership Report	139	0	0
Group Leave	0	0	0
Domain Wide Report	0	0	0
V3 Membership Report	136	0	0
Other Unknown types			0
IGMP v3 unsupported type			0
IGMP v3 source required for SSM			23
IGMP v3 mode not applicable for SSM			0

IGMP Global Statistics

Bad Length	0
Bad Checksum	0
Rx non-local	0

show igmp statistics

Syntax	show igmp statistics <brief detail> <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch and the QFX Series)	show igmp statistics <brief detail> <interface <i>interface-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display Internet Group Management Protocol (IGMP) statistics.
Options	none —Display IGMP statistics for all interfaces. brief detail —(Optional) Display the specified level of output. interface <i>interface-name</i> —(Optional) Display IGMP statistics about the specified interface only. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• clear igmp statistics on page 369
List of Sample Output	show igmp statistics on page 433 show igmp statistics interface on page 434
Output Fields	Table 38 on page 432 describes the output fields for the show igmp statistics command. Output fields are listed in the approximate order in which they appear.

Table 38: show igmp statistics Output Fields

Field Name	Field Description
IGMP packet statistics	Heading for IGMP packet statistics for all interfaces or for the specified interface name.

Table 38: show igmp statistics Output Fields (*continued*)

Field Name	Field Description
IGMP Message type	<p>Summary of IGMP statistics:</p> <ul style="list-style-type: none"> • Membership Query—Number of membership queries sent and received. • V1 Membership Report—Number of version 1 membership reports sent and received. • DVMRP—Number of DVMRP messages sent or received. • PIM V1—Number of PIM version 1 messages sent or received. • Cisco Trace—Number of Cisco trace messages sent or received. • V2 Membership Report—Number of version 2 membership reports sent or received. • Group Leave—Number of group leave messages sent or received. • Mtrace Response—Number of Mtrace response messages sent or received. • Mtrace Request—Number of Mtrace request messages sent or received. • Domain Wide Report—Number of domain-wide reports sent or received. • V3 Membership Report—Number of version 3 membership reports sent or received. • Other Unknown types—Number of unknown message types received. • IGMP v3 unsupported type—Number of messages received with unknown and unsupported IGMP version 3 message types. • IGMP v3 source required for SSM—Number of IGMP version 3 messages received that contained no source. • IGMP v3 mode not applicable for SSM—Number of IGMP version 3 messages received that did not contain a mode applicable for source-specific multicast (SSM).
Received	Number of messages received.
Sent	Number of messages sent.
Rx errors	Number of received packets that contained errors.
IGMP Global Statistics	<p>Summary of IGMP statistics for all interfaces.</p> <ul style="list-style-type: none"> • Bad Length—Number of messages received with length errors so severe that further classification could not occur. • Bad Checksum—Number of messages received with a bad IP checksum. No further classification was performed. • Bad Receive If—Number of messages received on an interface not enabled for IGMP. • Rx non-local—Number of messages received from senders that are not local. • Timed out—Number of groups that timed out as a result of not receiving an explicit leave message. • Rejected Report—Number of reports dropped because of the IGMP group policy. • Total Interfaces—Number of interfaces configured to support IGMP.

Sample Output

show igmp statistics

```

user@host> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type      Received      Sent  Rx errors
Membership Query        8883         459      0
V1 Membership Report     0            0      0

```

DVMRP	0	0	0
PIM V1	0	0	0
Cisco Trace	0	0	0
V2 Membership Report	0	0	0
Group Leave	0	0	0
Mtrace Response	0	0	0
Mtrace Request	0	0	0
Domain Wide Report	0	0	0
V3 Membership Report	0	0	0
Other Unknown types			0
IGMP v3 unsupported type			0
IGMP v3 source required for SSM			0
IGMP v3 mode not applicable for SSM			0
IGMP Global Statistics			
Bad Length	0		
Bad Checksum	0		
Bad Receive If	0		
Rx non-local	1227		
Timed out	0		
Rejected Report	0		
Total Interfaces	2		

show igmp statistics interface

```
user@host> show igmp statistics interface fe-1/0/1.0
IGMP interface packet statistics for fe-1/0/1.0
IGMP Message type      Received      Sent  Rx errors
Membership Query        0           230      0
V1 Membership Report    0           0        0
```

show mld group

Syntax	show mld group <brief detail> <group-name> <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display information about Multicast Listener Discovery (MLD) group membership.
Options	<p>none—Display standard information about all MLD groups.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>group-name—(Optional) Display MLD information about the specified group.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear mld membership on page 371
List of Sample Output	<p>show mld group (Include Mode) on page 436</p> <p>show mld group (Exclude Mode) on page 437</p> <p>show mld group brief on page 437</p> <p>show mld group detail (Include Mode) on page 437</p> <p>show mld group detail (Exclude Mode) on page 438</p>
Output Fields	Table 39 on page 435 describes the output fields for the show mld group command. Output fields are listed in the approximate order in which they appear.

Table 39: show mld group Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface that received the MLD membership report; local means that the local router joined the group itself.	All levels
Group	Group address.	All levels
Source	Source address.	All levels
Group Mode	Mode the SSM group is operating in: Include or Exclude .	All levels
Last reported by	Address of the host that last reported membership in this group.	All levels

Table 39: show mld group Output Fields (*continued*)

Field Name	Field Description	Level of Output
Source timeout	Time remaining until the group traffic is no longer forwarded. The timer is refreshed when a listener in include mode sends a report. A group in exclude mode or configured as a static group displays a zero timer.	detail
Timeout	Time remaining until the group membership is removed.	brief none
Group timeout	Time remaining until a group in exclude mode moves to include mode. The timer is refreshed when a listener in exclude mode sends a report. A group in include mode or configured as a static group displays a zero timer.	detail
Type	Type of group membership: <ul style="list-style-type: none"> • Dynamic—Host reported the membership. • Static—Membership is configured. 	All levels

Sample Output

show mld group (Include Mode)

```

user@host> show mld group
Interface: fe-0/1/2.0
  Group: ff02::1:ff05:1a67
    Group mode: Include
    Source: ::
    Last reported by: fe80::2e0:81ff:fe05:1a67
    Timeout:      245 Type: Dynamic
  Group: ff02::1:ffa8:c35e
    Group mode: Include
    Source: ::
    Last reported by: fe80::2e0:81ff:fe05:1a67
    Timeout:      241 Type: Dynamic
  Group: ff02::2:43e:d7f6
    Group mode: Include
    Source: ::
    Last reported by: fe80::2e0:81ff:fe05:1a67
    Timeout:      244 Type: Dynamic
  Group: ff05::2
    Group mode: Include
    Source: ::
    Last reported by: fe80::2e0:81ff:fe05:1a67
    Timeout:      244 Type: Dynamic
Interface: local
  Group: ff02::2
    Source: ::
    Last reported by: Local
    Timeout:      0 Type: Dynamic
  Group: ff02::16
    Source: ::
    Last reported by: Local
    Timeout:      0 Type: Dynamic

```

show mld group (Exclude Mode)

```

user@host> show mld group
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
  Group: ff02::6
    Source: ::
    Last reported by: fe80::21f:12ff:feb6:4b3a
    Timeout:      245 Type: Dynamic
  Group: ff02::16
    Source: ::
    Last reported by: fe80::21f:12ff:feb6:4b3a
    Timeout:      28 Type: Dynamic
Interface: local
  Group: ff02::2
    Source: ::
    Last reported by: Local
    Timeout:      0 Type: Dynamic
  Group: ff02::16
    Source: ::
    Last reported by: Local
    Timeout:      0 Type: Dynamic

```

show mld group brief

The output for the **show mld group brief** command is identical to that for the **show mld group** command. For sample output, see [show mld group \(Include Mode\) on page 436](#) and [show mld group \(Exclude Mode\) on page 437](#).

show mld group detail (Include Mode)

```

user@host> show mld group detail
Interface: fe-0/1/2.0
  Group: ff02::1:ff05:1a67
    Group mode: Include
    Source: ::
    Last reported by: fe80::2e0:81ff:fe05:1a67
    Timeout:      224 Type: Dynamic
  Group: ff02::1:ffa8:c35e
    Group mode: Include
    Source: ::
    Last reported by: fe80::2e0:81ff:fe05:1a67
    Timeout:      220 Type: Dynamic
  Group: ff02::2:43e:d7f6
    Group mode: Include
    Source: ::
    Last reported by: fe80::2e0:81ff:fe05:1a67
    Timeout:      223 Type: Dynamic
  Group: ff05::2
    Group mode: Include
    Source: ::
    Last reported by: fe80::2e0:81ff:fe05:1a67
    Timeout:      223 Type: Dynamic
Interface: so-1/0/1.0
  Group: ff02::2
    Group mode: Include
    Source: ::
    Last reported by: fe80::280:42ff:fe15:f445
    Timeout:      258 Type: Dynamic
Interface: local

```

```
Group: ff02::2
  Group mode: Include
  Source: ::
  Last reported by: Local
  Timeout:      0 Type: Dynamic
Group: ff02::16
  Source: ::
  Last reported by: Local
  Timeout:      0 Type: Dynamic
```

show mld group detail (Exclude Mode)

```
user@host> show mld group detail
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
  Group: ff02::6
    Group mode: Exclude
    Source: ::
    Source timeout: 0
    Last reported by: fe80::21f:12ff:feb6:4b3a
    Group timeout:   226 Type: Dynamic
  Group: ff02::16
    Group mode: Exclude
    Source: ::
    Source timeout: 0
    Last reported by: fe80::21f:12ff:feb6:4b3a
    Group timeout:   246 Type: Dynamic
Interface: local
  Group: ff02::2
    Group mode: Exclude
    Source: ::
    Source timeout: 0
    Last reported by: Local
    Group timeout:   0 Type: Dynamic
  Group: ff02::16
    Group mode: Exclude
    Source: ::
    Source timeout: 0
    Last reported by: Local
    Group timeout:   0 Type: Dynamic
```


show mld interface

Syntax	show mld interface <brief detail> <interface-name> <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display information about Multicast Listener Discovery (MLD)-enabled interfaces.
Options	<p>none—Display standard information about all MLD-enabled interfaces.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>interface-name—(Optional) Display information about the specified interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear mld membership on page 371
List of Sample Output	show mld interface on page 441 show mld interface brief on page 441 show mld interface detail on page 442 show mld interface <interface-name> on page 442
Output Fields	Table 40 on page 439 describes the output fields for the show mld interface command. Output fields are listed in the approximate order in which they appear.

Table 40: show mld interface Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface.	All levels
Querier	Address of the router that has been elected to send membership queries.	All levels
State	State of the interface: Up or Down .	All levels
SSM Map Policy	Name of the source-specific multicast (SSM) map policy that has been applied to the interface.	All levels
SSM Map Policy	Name of the source-specific multicast (SSM) map policy at the MLD interface.	All levels
Timeout	How long until the MLD querier is declared to be unreachable, in seconds.	All levels
Version	MLD version being used on the interface: 1 or 2.	All levels

Table 40: show mld interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
Groups	Number of groups on the interface.	All levels
Passive	<p>State of the passive mode option:</p> <ul style="list-style-type: none"> • On—Indicates that the router can run IGMP or MLD on the interface but not send or receive control traffic such as IGMP or MLD reports, queries, and leaves. • Off—Indicates that the router can run IGMP or MLD on the interface and send or receive control traffic such as IGMP or MLD reports, queries, and leaves. <p>The passive statement enables you to selectively activate up to two out of a possible three available query or control traffic options. When enabled, the following options appear after the on state declaration:</p> <ul style="list-style-type: none"> • send-general-query—The interface sends general queries. • send-group-query—The interface sends group-specific and group-source-specific queries. • allow-receive—The interface receives control traffic 	All levels
OIF map	Name of the OIF map associated to the interface.	All levels
SSM map	Name of the source-specific multicast (SSM) map used on the interface, if configured.	All levels
Group limit	Maximum number of groups allowed on the interface. Any memberships requested after the limit is reached are rejected.	All levels
Group threshold	<p>Configured threshold at which a warning message is generated.</p> <p>This threshold is based on a percentage of groups received on the interface. If the number of groups received reaches the configured threshold, the device generates a warning message.</p>	All levels
Group log-interval	Time (in seconds) between consecutive log messages.	All levels
Immediate Leave	<p>State of the immediate leave option:</p> <ul style="list-style-type: none"> • On—Indicates that the router removes a host from the multicast group as soon as the router receives a multicast listener done message from a host associated with the interface. • Off—Indicates that after receiving a multicast listener done message, instead of removing a host from the multicast group immediately, the router sends a group query to determine if another receiver responds. 	All levels

Table 40: show mld interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
Configured Parameters	<p>Information configured by the user.</p> <ul style="list-style-type: none"> • MLD Query Interval (.1 secs)—Interval at which this router sends membership queries when it is the querier. • MLD Query Response Interval (.1 secs)—Time that the router waits for a report in response to a general query. • MLD Last Member Query Interval (.1 secs)—Time that the router waits for a report in response to a group-specific query. • MLD Robustness Count—Number of times the router retries a query. 	All levels
Derived Parameters	<p>Derived information.</p> <ul style="list-style-type: none"> • MLD Membership Timeout (.1 secs)—Timeout period for group membership. If no report is received for these groups before the timeout expires, the group membership will be removed. • MLD Other Querier Present Timeout (.1 secs)—Time that the router waits for the IGMP querier to send a query. 	All levels

Sample Output

show mld interface

```

user@host> show mld interface
Interface: fe-0/0/0
  Querier: None
  State: Up      Timeout:      0   Version:  1   Groups:    0
  SSM Map Policy: ssm-policy-A
Interface: at-0/3/1.0
  Querier: 8038::c0a8:c345
  State: Up      Timeout:    None   Version:  1   Groups:    0
  SSM Map Policy: ssm-policy-B
Interface: fe-1/0/1.0
  Querier: ::192.168.195.73
  State: Up      Timeout:    None   Version:  1   Groups:    3
  SSM Map Policy: ssm-policy-C
  SSM map: ipv6map1
Immediate Leave: On

Configured Parameters:
MLD Query Interval (.1 secs): 1250
MLD Query Response Interval (.1 secs): 100
MLD Last Member Query Interval (.1 secs): 10
MLD Robustness Count: 2

Derived Parameters:
MLD Membership Timeout (.1secs): 2600
MLD Other Querier Present Timeout (.1 secs): 2550

```

show mld interface brief

The output for the **show mld interface brief** command is identical to that for the **show mld interface** command. For sample output, see [show mld interface on page 441](#).

show mld interface detail

The output for the **show mld interface detail** command is identical to that for the **show mld interface** command. For sample output, see [show mld interface on page 441](#).

show mld interface <interface-name>

```
user@host# show mld interface ge-3/2/0.0
Interface: ge-3/2/0.0
Querier: 20.1.1.1
State: Up Timeout:      None Version: 3 Groups:      1
Group limit: 8
Group threshold: 60
Group log-interval: 10
Immediate leave: Off
Promiscuous mode: Off
```

show mld statistics

Syntax	show mld statistics <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display information about Multicast Listener Discovery (MLD) statistics.
Options	<p>none—Display MLD statistics for all interfaces.</p> <p>interface <i>interface-name</i>—(Optional) Display statistics about the specified interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear mld statistics on page 372
List of Sample Output	show mld statistics on page 444 show mld statistics interface on page 445
Output Fields	<p>Table 41 on page 443 describes the output fields for the show mld statistics command. Output fields are listed in the approximate order in which they appear.</p>

Table 41: show mld statistics Output Fields

Field Name	Field Description
Received	Number of received packets.
Sent	Number of transmitted packets.
Rx errors	Number of received packets that contained errors.

Table 41: show mld statistics Output Fields (*continued*)

Field Name	Field Description
MLD Message type	Summary of MLD statistics. <ul style="list-style-type: none"> • Listener Query (v1/v2)—Number of membership queries sent and received. • Listener Report (v1)—Number of version 1 membership reports sent and received. • Listener Done (v1/v2)—Number of Listener Done messages sent and received. • Listener Report (v2)—Number of version 2 membership reports sent and received. • Other Unknown types—Number of unknown message types received. • MLD v2 source required for SSM—Number of MLD version 2 messages received that contained no source. • MLD v2 mode not applicable for SSM—Number of MLD version 2 messages received that did not contain a mode applicable for source-specific multicast (SSM).
MLD Global Statistics	Summary of MLD statistics for all interfaces. <ul style="list-style-type: none"> • Bad Length—Number of messages received with length errors so severe that further classification could not occur. • Bad Checksum—Number of messages received with an invalid IP checksum. No further classification was performed. • Bad Receive If—Number of messages received on an interface not enabled for MLD. • Rx non-local—Number of messages received from nonlocal senders. • Timed out—Number of groups that timed out as a result of not receiving an explicit leave message. • Rejected Report—Number of reports dropped because of the MLD group policy. • Total Interfaces—Number of interfaces configured to support IGMP.

Sample Output

show mld statistics

```

user@host> show mld statistics
MLD packet statistics for all interfaces
MLD Message type      Received      Sent  Rx errors
Listener Query (v1/v2)    0            2      0
Listener Report (v1)      0            0      0
Listener Done (v1/v2)     0            0      0
Listener Report (v2)      0            0      0
Other Unknown types      0            0      0
MLD v2 source required for SSM  2
MLD v2 mode not applicable for SSM 0

MLD Global Statistics
Bad Length              0
Bad Checksum            0
Bad Receive If          0
Rx non-local            0
Timed out               0

```

Rejected Report	0
Total Interfaces	2

show mld statistics interface

```
user@host> show mld statistics interface fe-1/0/1.0
MLD interface packet statistics for fe-1/0/1.0
MLD Message type      Received      Sent  Rx errors
Listener Query (v1/v2)    0           2      0
Listener Report (v1)      0           0      0
Listener Done (v1/v2)     0           0      0
Listener Report (v2)      0           0      0
Other Unknown types              0      0
MLD v2 source required for SSM    2
MLD v2 mode not applicable for SSM 0

MLD Global Statistics
Bad Length                0
Bad Checksum              0
Bad Receive If            0
Rx non-local              0
Timed out                 0
Rejected Report           0
Total Interfaces          2
```

show msdp

Syntax	<pre>show msdp <brief detail> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <peer <i>peer-address</i>></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Display Multicast Source Discovery Protocol (MSDP) information.
Options	<p>none—Display standard MSDP information for all routing instances.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display information for the specified instance only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>peer <i>peer-address</i>—(Optional) Display information about the specified peer only.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show msdp source on page 448 • show msdp source-active on page 450 • show msdp statistics on page 453
List of Sample Output	<p>show msdp on page 447</p> <p>show msdp brief on page 447</p> <p>show msdp detail on page 447</p>
Output Fields	Table 42 on page 446 describes the output fields for the show msdp command. Output fields are listed in the approximate order in which they appear.

Table 42: show msdp Output Fields

Field Name	Field Description	Level of Output
Peer address	IP address of the peer.	All levels
Local address	Local address of the peer.	All levels
State	Status of the MSDP connection: Listen , Established , or Inactive .	All levels
Last up/down	Time at which the most recent peer-state change occurred.	All levels

Table 42: show msdp Output Fields (*continued*)

Field Name	Field Description	Level of Output
Peer-Group	Peer group name.	All levels
SA Count	Number of source-active cache entries advertised by each peer that were accepted, compared to the number that were received, in the format <i>number-accepted/number-received</i> .	All levels
Peer Connect Retries	Number of peer connection retries.	detail
State timer expires	Number of seconds before another message is sent to a peer.	detail
Peer Times out	Number of seconds to wait for a response from the peer before the peer is declared unavailable.	detail
SA accepted	Number of entries in the source-active cache accepted from the peer.	detail
SA received	Number of entries in the source-active cache received by the peer.	detail

Sample Output

show msdp

```

user@host> show msdp
Peer address    Local address  State      Last up/down Peer-Group SA Count
198.32.8.193    198.32.8.195  Established 5d 19:25:44 North23 120/150
198.32.8.194    198.32.8.195  Established 3d 19:27:27 North23 300/345
198.32.8.196    198.32.8.195  Established 5d 19:39:36 North23 10/13
198.32.8.197    198.32.8.195  Established 5d 19:32:27 North23 5/6
198.32.8.198    198.32.8.195  Established 3d 19:33:04 North23 2305/3000

```

show msdp brief

The output for the **show msdp brief** command is identical to that for the **show msdp** command. For sample output, see [show msdp on page 447](#).

show msdp detail

```

user@host> show msdp detail
Peer: 10.255.70.15
Local address: 10.255.70.19
State: Established
Peer Connect Retries: 0
State timer expires: 22
Peer Times out: 49
SA accepted: 0
SA received: 0

```

show msdp source

Syntax	<code>show msdp source</code> <code><instance <i>instance-name</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code> <code><source-address></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display multicast sources learned from Multicast Source Discovery Protocol (MSDP).
Options	none —Display standard MSDP source information for all routing instances. instance <i>instance-name</i> —(Optional) Display information for the specified instance only. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. source-address —(Optional) IP address and optional prefix length. Display information for the specified source address only.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show msdp on page 446• show msdp source-active on page 450• show msdp statistics on page 453
List of Sample Output	show msdp source on page 449

Output Fields Table 43 on page 449 describes the output fields for the **show msdp source** command. Output fields are listed in the approximate order in which they appear.

Table 43: show msdp source Output Fields

Field Name	Field Description
Source address	IP address of the source.
/Len	Length of the prefix for this IP address.
Type	Discovery method for this multicast source: <ul style="list-style-type: none"> • Configured—Source-active limit explicitly configured for this source. • Dynamic—Source-active limit established when this source was discovered.
Maximum	Source-active limit applied to this source.
Threshold	Source-active threshold applied to this source.
Exceeded	Number of source-active messages received from this source exceeding the established maximum.

Sample Output

show msdp source

```

user@host> show msdp source
Source address /Len  Type      Maximum  Threshold  Exceeded
0.0.0.0       /0    Configured    5         none        0
10.1.0.0      /16   Configured    500        none        0
10.1.1.1      /32   Configured   10000       none        0
10.1.1.2      /32   Dynamic     6936        none        0
10.1.5.5      /32   Dynamic     500         none       123
10.2.1.1      /32   Dynamic      2           none        0

```

show msdp source-active

Syntax	<code>show msdp source-active</code> <code><brief detail></code> <code><group <i>group</i>></code> <code><instance <i>instance-name</i>></code> <code><local></code> <code><logical-system (all <i>logical-system-name</i>)></code> <code><originator <i>originator</i>></code> <code><peer <i>peer-address</i>></code> <code><source <i>source-address</i>></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display the Multicast Source Discovery Protocol (MSDP) source-active cache.
Options	none —Display standard MSDP source-active cache information for all routing instances. brief detail —(Optional) Display the specified level of output. group <i>group</i> —(Optional) Display source-active cache information for the specified group. instance <i>instance-name</i> —(Optional) Display information for the specified instance. local —(Optional) Display all source-active caches originated by this router. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. originator <i>originator</i> —(Optional) Display information about the peer that originated the source-active cache entries. peer <i>peer-address</i> —(Optional) Display the source-active cache of the specified peer. source <i>source-address</i> —(Optional) Display the source-active cache of the specified source.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show msdp on page 446• show msdp source on page 448• show msdp statistics on page 453
List of Sample Output	show msdp source-active on page 451 show msdp source-active brief on page 451 show msdp source-active detail on page 452 show msdp source-active source on page 452
Output Fields	Table 44 on page 451 describes the output fields for the show msdp source-active command. Output fields are listed in the approximate order in which they appear.

Table 44: show msdp source-active Output Fields

Field Name	Field Description
Global active source limit exceeded	Number of times all peers have exceeded configured active source limits.
Global active source limit maximum	Configured number of active source messages accepted by the device.
Global active source limit threshold	Configured threshold for applying random early discard (RED) to drop some but not all MSDP active source messages.
Global active source limit log-warning	Threshold at which a warning message is logged (percentage of the number of active source messages accepted by the device).
Global active source limit log interval	Time (in seconds) between consecutive log messages.
Group address	Multicast address of the group.
Source address	IP address of the source.
Peer address	IP address of the peer.
Originator	Router ID configured on the source of the rendezvous point (RP) that originated the message, or the loopback address when the router ID is not configured.
Flags	Flags: Accept, Reject, or Filtered.

Sample Output

show msdp source-active

```

user@host> show msdp source-active
Group address  Source address  Peer address  Originator  Flags
230.0.0.0      192.168.195.46  local        10.255.14.30  Accept
230.0.0.1      192.168.195.46  local        10.255.14.30  Accept
230.0.0.2      192.168.195.46  local        10.255.14.30  Accept
230.0.0.3      192.168.195.46  local        10.255.14.30  Accept
230.0.0.4      192.168.195.46  local        10.255.14.30  Accept

```

show msdp source-active brief

The output for the **show msdp source-active brief** command is identical to that for the **show msdp source-active** command. For sample output, see [show msdp source-active on page 451](#).

show msdp source-active detail

The output for the **show msdp source-active detail** command is identical to that for the **show msdp source-active** command. For sample output, see [show msdp source-active on page 451](#).

show msdp source-active source

```
user@host> show msdp source-active source 192.168.215.246
```

```
Global active source limit exceeded: 0
```

```
Global active source limit maximum: 25000
```

```
Global active source limit threshold: 24000
```

```
Global active source limit log-warning: 100
```

```
Global active source limit log interval: 0
```

Group address	Source address	Peer address	Originator	Flags
226.2.2.1	192.168.215.246	10.255.182.140	10.255.182.140	Accept
226.2.2.3	192.168.215.246	10.255.182.140	10.255.182.140	Accept
226.2.2.4	192.168.215.246	10.255.182.140	10.255.182.140	Accept
226.2.2.5	192.168.215.246	10.255.182.140	10.255.182.140	Accept
226.2.2.7	192.168.215.246	10.255.182.140	10.255.182.140	Accept
226.2.2.10	192.168.215.246	10.255.182.140	10.255.182.140	Accept
226.2.2.11	192.168.215.246	10.255.182.140	10.255.182.140	Accept
226.2.2.13	192.168.215.246	10.255.182.140	10.255.182.140	Accept
226.2.2.14	192.168.215.246	10.255.182.140	10.255.182.140	Accept
226.2.2.15	192.168.215.246	10.255.182.140	10.255.182.140	Accept

show msdp statistics

Syntax	show msdp statistics <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <peer <i>peer-address</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display statistics about Multicast Source Discovery Protocol (MSDP) peers.
Options	<p>none—Display statistics about all MSDP peers for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Display statistics about a specific MSDP instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>peer <i>peer-address</i>—(Optional) Display statistics about a particular MSDP peer.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear msdp statistics on page 374
List of Sample Output	show msdp statistics on page 455 show msdp statistics peer on page 455
Output Fields	Table 45 on page 453 describes the output fields for the show msdp statistics command. Output fields are listed in the approximate order in which they appear.

Table 45: show msdp statistics Output Fields

Field Name	Field Description
Global active source limit exceeded	Number of times all peers have exceeded configured active source limits.
Global active source limit maximum	Configured number of active source messages accepted by the device.
Global active source limit threshold	Configured threshold for applying random early discard (RED) to drop some but not all MSDP active source messages.
Global active source limit log-warning	Threshold at which a warning message is logged (percentage of the number of active source messages accepted by the device).
Global active source limit log interval	Time (in seconds) between consecutive log messages.
Peer	Address of peer.

Table 45: show msdp statistics Output Fields (*continued*)

Field Name	Field Description
Last State Change	How long ago the peer state changed.
Last message received from the peer	How long ago the last message was received from the peer.
RPF Failures	Number of reverse path forwarding (RPF) failures.
Remote Closes	Number of times the remote peer closed.
Peer Timeouts	Number of peer timeouts.
SA messages sent	Number of source-active messages sent.
SA messages received	Number of source-active messages received.
SA request messages sent	Number of source-active request messages sent.
SA request messages received	Number of source-active request messages received.
SA response messages sent	Number of source-active response messages sent.
SA response messages received	Number of source-active response messages received.
Active source exceeded	Number of times this peer has exceeded configured source-active limits.
Active source Maximum	Configured number of active source messages accepted by this peer.
Active source threshold	Configured threshold on this peer for applying random early discard (RED) to drop some but not all MSDP active source messages.
Active source log-warning	Configured threshold on this peer at which a warning message is logged (percentage of the number of active source messages accepted by the device).
Active source log-interval	Time (in seconds) between consecutive log messages on this peer.
Keepalive messages sent	Number of keepalive messages sent.
Keepalive messages received	Number of keepalive messages received.
Unknown messages received	Number of unknown messages received.

Table 45: show msdp statistics Output Fields (*continued*)

Field Name	Field Description
Error messages received	Number of error messages received.

Sample Output

show msdp statistics

```

user@host> show msdp statistics
Global active source limit exceeded: 0
Global active source limit maximum: 10
Global active source limit threshold: 8
Global active source limit log-warning: 60
Global active source limit log interval: 60

Peer: 10.255.245.39
Last State Change: 11:54:49 (00:24:59)
Last message received from peer: 11:53:32 (00:26:16)
RPF Failures: 0
Remote Closes: 0
Peer Timeouts: 0
SA messages sent: 376
SA messages received: 459
SA request messages sent: 0
SA request messages received: 0
SA response messages sent: 0
SA response messages received: 0
Active source exceeded: 0
Active source Maximum: 10
Active source threshold: 8
Active source log-warning: 60
Active source log-interval 120
Keepalive messages sent: 17
Keepalive messages received: 19
Unknown messages received: 0
Error messages received: 0

```

show msdp statistics peer

```

user@host> show msdp statistics peer 10.255.182.140
Peer: 10.255.182.140
  Last State Change: 8:19:23 (00:01:08)
  Last message received from peer: 8:20:05 (00:00:26)
  RPF Failures: 0
  Remote Closes: 0
  Peer Timeouts: 0
  SA messages sent: 17
  SA messages received: 16
  SA request messages sent: 0
  SA request messages received: 0
  SA response messages sent: 0
  SA response messages received: 0
  Active source exceeded: 20
  Active source Maximum: 10
  Active source threshold: 8
  Active source log-warning: 60
  Active source log-interval: 120
  Keepalive messages sent: 0

```

Keepalive messages received: 0
Unknown messages received: 0
Error messages received: 0

show multicast backup-pe-groups

Syntax	show multicast backup-pe-groups <address <i>pe-address</i> > <group <i>group-name</i> > <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced in Junos OS Release 9.0.
Description	Display backup PE router group information when ingress PE redundancy is configured. Ingress PE redundancy provides a backup resource when point-to-multipoint LSPs are configured for multicast distribution.
Options	<p>none—Display standard information about all backup PE groups.</p> <p>address <i>pe-address</i>—(Optional) Display the groups that a PE address is associated with.</p> <p>group <i>group</i>—(Optional) Display the backup PE group information for a particular group.</p> <p>instance <i>instance-name</i>—(Optional) Display backup PE group information for a specific multicast instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show multicast backup-pe-groups on page 458
Output Fields	Table 46 on page 457 describes the output fields for the show multicast backup-pe-groups command. Output fields are listed in the approximate order in which they appear.

Table 46: show multicast backup-pe-groups Output Fields

Field Name	Field Description
Backup PE Group	Group name.
Designated PE	Primary PE router. Address of the PE router that is currently forwarding traffic on the static route.
Transitions	Number of times that the designated PE router has transitioned from the most eligible PE router to a backup PE router and back again to the most eligible PE router.
Last Transition	Time of the most recent transition.
Local Address	Address of the local PE router.
Backup PE List	List of PE routers that are configured to be backups for the group.

Sample Output

show multicast backup-pe-groups

```
user@host> show multicast backup-pe-groups
Instance: master

Backup PE group: b1
  Designated PE: 10.255.165.7
  Transitions: 1
  Last Transition: 03:15:01
  Local Address: 10.255.165.7
  Backup PE List:
    10.255.165.8

Backup PE group: b2
  Designated PE: 10.255.165.7
  Transitions: 2
  Last Transition: 02:58:20
  Local Address: 10.255.165.7
  Backup PE List:
    10.255.165.9
    10.255.165.8
```

show multicast flow-map

Syntax	show multicast flow-map <brief detail> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch and the QFX Series)	show multicast flow-map <brief detail>
Release Information	Command introduced in Junos OS Release 8.2. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display configuration information about IP multicast flow maps.
Options	none —Display configuration information about IP multicast flow maps on all systems. brief detail —(Optional) Display the specified level of output. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show multicast flow-map on page 460 show multicast flow-map detail on page 460
Output Fields	Table 47 on page 459 describes the output fields for the show multicast flow-map command. Output fields are listed in the approximate order in which they appear.

Table 47: show multicast flow-map Output Fields

Field Name	Field Description	Levels of Output
Name	Name of the flow map.	All levels
Policy	Name of the policy associated with the flow map.	All levels
Cache-timeout	Cache timeout value assigned to the flow map.	All levels
Bandwidth	Bandwidth setting associated with the flow map.	All levels
Adaptive	Whether or not adaptive mode is enabled for the flow map.	none
Flow-map	Name of the flow map.	detail
Adaptive Bandwidth	Whether or not adaptive mode is enabled for the flow map.	detail
Redundant Sources	Redundant sources defined for the same destination group.	detail

Sample Output

show multicast flow-map

```
user@host> show multicast flow-map
Instance: master
Name      Policy      Cache timeout      Bandwidth Adaptive
map2      policy2     never              2000000 no
map1      policy1     60 seconds        2000000 no
```

Sample Output

show multicast flow-map detail

```
user@host> show multicast flow-map detail
Instance: master
Flow-map: map1
  Policy:      policy1
  Cache Timeout: 600 seconds
  Bandwidth:   2000000
  Adaptive Bandwidth: yes
  Redundant Sources: 11.11.11.11
  Redundant Sources: 11.11.11.12
  Redundant Sources: 11.11.11.13
```

show multicast forwarding-cache statistics

Syntax	show multicast forwarding-cache statistics <inet inet6> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced in Junos OS Release 12.2.
Description	Display IP multicast forwarding cache statistics.
Options	<p>none—Display multicast forwarding cache statistics for all supported address families for all routing instances.</p> <p>inet inet6—(Optional) Display multicast forwarding cache statistics for IPv4 or IPv6 family addresses, respectively.</p> <p>instance <i>instance-name</i>—(Optional) Display multicast forwarding cache statistics for a specific routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear multicast forwarding-cache on page 377
List of Sample Output	show multicast forwarding-cache statistics on page 462 show multicast forwarding-cache statistics instance on page 462
Output Fields	Table 48 on page 461 describes the output fields for the show multicast forwarding-cache statistics command. Output fields are listed in the approximate order in which they appear.

Table 48: show multicast forwarding-cache statistics Output Fields

Field Name	Field Description
Instance	Name of the routing instance for which multicast forwarding cache statistics are displayed.
Family	Protocol family for which multicast forwarding cache statistics are displayed: ALL , INET , or INET6 .
Suppress Threshold	Maximum number of multicast forwarding cache entries that can be added to the cache. When the number of entries reaches the configured threshold, the device suspends adding new multicast forwarding cache entries.
Reuse Value	Number of multicast forwarding cache entries that must be reached before the device creates new multicast forwarding cache entries. When the total number of multicast forwarding cache entries is below the reuse value, the device resumes adding new multicast forwarding cache entries.

Table 48: show multicast forwarding-cache statistics Output Fields (*continued*)

Field Name	Field Description
Warning Threshold	Threshold at which a warning message is logged (percentage of the suppress threshold).
Currently Used Entries	Number of currently used multicast forwarding cache entries.

Sample Output

show multicast forwarding-cache statistics

```
user@host> show multicast forwarding-cache statistics
Instance: master Family: INET
Suppress Threshold           100
Reuse Value                  80
Warning Threshold           90
Currently Used Entries       101

Instance: master Family: INET6
Suppress Threshold           50
Reuse Value                  50
Warning Threshold           80
Currently Used Entries        3
```

show multicast forwarding-cache statistics instance

```
user@host> show multicast forwarding-cache statistics instance VPN-A
Instance: VPN-A Family: ALL
Suppress Threshold           20
Reuse Value                  16
Warning Threshold           50
Currently Used Entries       17
```


show multicast interface

Syntax	show multicast interface <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch and the QFX Series)	show multicast interface
Release Information	Command introduced in Junos OS Release 8.3. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display bandwidth information about IP multicast interfaces.
Options	none —Display all interfaces that have multicast configured. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show multicast interface on page 464
Output Fields	Table 49 on page 463 describes the output fields for the show multicast interface command. Output fields are listed in the approximate order in which they appear.

Table 49: show multicast interface Output Fields

Field Name	Field Description
Interface	Name of the multicast interface.
Maximum bandwidth (bps)	Maximum bandwidth setting, in bits per second, for this interface.
Remaining bandwidth (bps)	Amount of bandwidth, in bits per second, remaining on the interface.
Mapped bandwidth deduction (bps)	Amount of bandwidth, in bits per second, used by any flows that are mapped to the interface. NOTE: Adding the mapped bandwidth deduction value to the local bandwidth deduction value results in the total deduction value for the interface. This field does not appear in the output when the no QoS adjustment feature is disabled.

Table 49: show multicast interface Output Fields (*continued*)

Field Name	Field Description
Local bandwidth deduction (bps)	<p>Amount of bandwidth, in bits per second, used by any mapped flows that are traversing the interface.</p> <p>NOTE: Adding the mapped bandwidth deduction value to the local bandwidth deduction value results in the total deduction value for the interface.</p> <p>This field does not appear in the output when the no QoS adjustment feature is disabled.</p>
Reverse OIF mapping	<p>State of the reverse OIF mapping feature (on or off).</p> <p>NOTE: This field does not appear in the output when the no QoS adjustment feature is disabled.</p>
Reverse OIF mapping no QoS adjustment	<p>State of the no QoS adjustment feature (on or off) for interfaces that are using reverse OIF mapping.</p> <p>NOTE: This field does not appear in the output when the no QoS adjustment feature is disabled.</p>
Leave timer	<p>Amount of time a mapped interface remains active after the last mapping ends.</p> <p>NOTE: This field does not appear in the output when the no QoS adjustment feature is disabled.</p>
No QoS adjustment	<p>State (on) of the no QoS adjustment feature when this feature is enabled.</p> <p>NOTE: This field does not appear in the output when the no QoS adjustment feature is disabled.</p>

Sample Output

show multicast interface

```

user@host> show multicast interface
Interface           Maximum bandwidth (bps) Remaining bandwidth (bps)
fe-0/0/3            10000000                0
fe-0/0/3.210        10000000                -2000000
fe-0/0/3.220        100000000               100000000
fe-0/0/3.230        20000000                18000000
fe-0/0/2.200        100000000               100000000

```

show multicast mrimfo

Syntax	<code>show multicast mrimfo</code> <code><host></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display configuration information about IP multicast networks, including neighboring multicast router addresses.
Options	none —Display configuration information about all multicast networks. host —(Optional) Display configuration information about a particular host. Replace <i>host</i> with a hostname or IP address.
Required Privilege Level	view
List of Sample Output	show multicast mrimfo on page 466
Output Fields	Table 50 on page 465 describes the output fields for the show multicast mrimfo command. Output fields are listed in the approximate order in which they appear.

Table 50: show multicast mrimfo Output Fields

Field Name	Field Description
<i>source-address</i>	Query address, hostname (DNS name or IP address of the source address), and multicast protocol version or the software version of another vendor.
<i>ip-address-1—>ip-address-2</i>	Queried router interface address and directly attached neighbor interface address, respectively.
<i>(name or ip-address)</i>	Name or IP address of neighbor.
<i>[metric/threshold/type/flags]</i>	Neighbor's multicast profile: <ul style="list-style-type: none"> metric—Always has a value of 1, because mrimfo queries the directly connected interfaces of a device. threshold—Multicast threshold time-to-live (TTL). The range of values is 0 through 255. type—Multicast connection type: pim or tunnel. flags—Flags for this route: <ul style="list-style-type: none"> querier—Queried router is the designated router for the neighboring session. leaf—Link is a leaf in the multicast network. down—Link status indicator.

Sample Output

show multicast mrinfo

```
user@host> show multicast mrinfo 10.35.4.1
10.35.4.1 (10.35.4.1) [version 12.0]:
  192.168.195.166 -> 0.0.0.0 (local) [1/0/pim/querier/leaf]
  10.38.20.1 -> 0.0.0.0 (local) [1/0/pim/querier/leaf]
  10.47.1.1 -> 10.47.1.2 (10.47.1.2) [1/5/pim]
  0.0.0.0 -> 0.0.0.0 (local) [1/0/pim/down]
```

show multicast next-hops

Syntax	<pre>show multicast next-hops <brief detail> <identifier-number> <inet inet6> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switch and the QFX Series)	<pre>show multicast next-hops <brief detail> <identifier-number> <inet inet6></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>inet6 option introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>detail option display of next-hop ID number introduced in Junos OS Release 11.1 for M Series and T Series routers and EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p>
Description	Display the entries in the IP multicast next-hop table.
Options	<p>none—Display standard information about all entries in the multicast next-hop table for all supported address families.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>When you include the detail option on M Series and T Series routers and EX Series switches, the downstream interface name includes the next-hop ID number in parentheses, in the form fe-0/1/2.0-(1048574) where 1048574 is the next-hop ID number.</p> <p>identifier-number—(Optional) Show a particular next hop by ID number. The range of values is 1 through 65,535.</p> <p>inet inet6—(Optional) Display entries for IPv4 or IPv6 family addresses, respectively.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	<p>show multicast next-hops on page 468</p> <p>show multicast next-hops (Bidirectional PIM on page 468)</p> <p>show multicast next-hops brief on page 469</p> <p>show multicast next-hops detail on page 469</p>
Output Fields	<p>Table 51 on page 468 describes the output fields for the show multicast next-hops command. Output fields are listed in the approximate order in which they appear.</p>

Table 51: show multicast next-hops Output Fields

Field Name	Field Description
Family	Protocol family (such as INET).
ID	Next-hop identifier of the prefix. The identifier is returned by the routing device's Packet Forwarding Engine.
RefCount	Number of cache entries that are using this next hop.
KRefCount	Kernel reference count for the next hop.
Downstream interface	Interface names associated with each multicast next-hop ID.
Incoming interface list	List of interfaces that accept incoming traffic. Only shown for routes that do not use strict RPF-based forwarding, for example for bidirectional PIM.

Sample Output

show multicast next-hops

```

user@host> show multicast next-hops
Family: INET
ID      Refcount  KRefCount Downstream interface
262142      4          2 so-1/0/0.0
262143      2          1 mt-1/1/0.49152
262148      2          1 mt-1/1/0.32769

```

show multicast next-hops (Bidirectional PIM)

```

user@host> show multicast next-hops
Family: INET
ID      Refcount  KRefCount Downstream interface
2097151      8          4 ge-0/0/1.0

Family: INET6
ID      Refcount  KRefCount Downstream interface
2097157      2          1 ge-0/0/1.0

Family: Incoming interface list
ID      Refcount  KRefCount Downstream interface
513      5          2 lo0.0
           ge-0/0/1.0
514      5          2 lo0.0
           ge-0/0/1.0
           xe-4/1/0.0
515      3          1 lo0.0
           ge-0/0/1.0
           xe-4/1/0.0
544      1          0 lo0.0
           xe-4/1/0.0

```

show multicast next-hops brief

The output for the **show multicast next-hops brief** command is identical to that for the **show multicast next-hops** command. For sample output, see [show multicast next-hops on page 468](#).

show multicast next-hops detail

```
user@host> show multicast next-hops detail
Family: INET
ID          Refcount KRefCount Downstream interface
1048577      2          1 fe-0/1/2.0-(1048574)
              ge-0/2/3.0-(1048576)
```

show multicast pim-to-igmp-proxy

Syntax	show multicast pim-to-igmp-proxy <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch and the QFX Series)	show multicast pim-to-igmp-proxy <instance <i>instance-name</i> >
Release Information	Command introduced in Junos OS Release 9.6. Command introduced in Junos OS Release 9.6 for EX Series switches. instance option introduced in Junos OS Release 10.3. instance option introduced in Junos OS Release 10.3 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display configuration information about PIM-to-IGMP message translation, also known as PIM-to-IGMP proxy.
Options	none —Display configuration information about PIM-to-IGMP message translation for all routing instances. instance <i>instance-name</i> —(Optional) Display configuration information about PIM-to-IGMP message translation for a specific multicast instance. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring PIM-to-IGMP and PIM-to-MLD Message Translation</i>
List of Sample Output	show multicast pim-to-igmp-proxy on page 471 show multicast pim-to-igmp-proxy instance on page 471
Output Fields	Table 52 on page 470 describes the output fields for the show multicast pim-to-igmp-proxy command. Output fields are listed in the order in which they appear.

Table 52: show multicast pim-to-igmp-proxy Output Fields

Field Name	Field Description
Instance	Routing instance. Default instance is master (inet.0 routing table).
Proxy state	State of PIM-to-IGMP message translation, also known as PIM-to-IGMP proxy, on the configured upstream interfaces: enabled or disabled .
<i>interface-name</i>	Name of upstream interface (no more than two allowed) on which PIM-to-IGMP message translation is configured.

Sample Output

show multicast pim-to-igmp-proxy

```
user@host> show multicast pim-to-igmp-proxy
Instance: master Proxy state: enabled
ge-0/1/0.1
ge-0/1/0.2
```

show multicast pim-to-igmp-proxy instance

```
user@host> show multicast pim-to-igmp-proxy instance VPN-A
Instance: VPN-A Proxy state: enabled
ge-0/1/0.1
```

show multicast pim-to-mld-proxy

Syntax	show multicast pim-to-mld-proxy <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch and the QFX Series)	show multicast pim-to-mld-proxy <instance <i>instance-name</i> >
Release Information	Command introduced in Junos OS Release 9.6. Command introduced in Junos OS Release 9.6 for EX Series switches. instance option introduced in Junos OS Release 10.3. instance option introduced in Junos OS Release 10.3 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display configuration information about PIM-to-MLD message translation, also known as PIM-to-MLD proxy.
Options	none —Display configuration information about PIM-to-MLD message translation for all routing instances. instance <i>instance-name</i> —(Optional) Display configuration information about PIM-to-MLD message translation for a specific multicast instance. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show multicast pim-to-mld-proxy on page 473 show multicast pim-to-mld-proxy instance on page 473
Output Fields	Table 53 on page 472 describes the output fields for the show multicast pim-to-mld-proxy command. Output fields are listed in the order in which they appear.

Table 53: show multicast pim-to-mld-proxy Output Fields

Field Name	Field Description
Proxy state	State of PIM-to-MLD message translation, also known as PIM-to-MLD proxy, on the configured upstream interfaces: enabled or disabled .
<i>interface-name</i>	Name of upstream interface (no more than two allowed) on which PIM-to-MLD message translation is configured.

Sample Output

show multicast pim-to-mld-proxy

```
user@host> show multicast pim-to-mld-proxy
Instance: master Proxy state: enabled
ge-0/5/0.1
ge-0/5/0.2
```

show multicast pim-to-mld-proxy instance

```
user@host> show multicast pim-to-mld-proxy instance VPN-A
Instance: VPN-A Proxy state: enabled
ge-0/5/0.1
```

show multicast route

Syntax	<pre>show multicast route <brief detail extensive summary> <active all inactive> <group group> <inet inet6> <instance instance name> <logical-system (all logical-system-name)> <regular-expression> <source-prefix source-prefix></pre>
Syntax (EX Series Switch and the QFX Series)	<pre>show multicast route <brief detail extensive summary> <active all inactive> <group group> <inet inet6> <instance instance name> <regular-expression> <source-prefix source-prefix></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p>
Description	Display the entries in the IP multicast forwarding table. You can display similar information with the show route table inet.1 command.
Options	<p>none—Display standard information about all entries in the multicast forwarding table for all routing instances.</p> <p>brief detail extensive summary—(Optional) Display the specified level of output.</p> <p>active all inactive—(Optional) Display all active entries, all entries, or all inactive entries, respectively, in the multicast forwarding table.</p> <p>group group—(Optional) Display the cache entries for a particular group.</p> <p>inet inet6—(Optional) Display multicast forwarding table entries for IPv4 or IPv6 family addresses, respectively.</p> <p>instance instance-name—(Optional) Display entries in the multicast forwarding table for a specific multicast instance.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>regular-expression—(Optional) Display information about the multicast forwarding table entries that match a UNIX OS-style regular expression.</p>

source-prefix *source-prefix*—(Optional) Display the cache entries for a particular source prefix.

Required Privilege Level view

Related Documentation • *Example: Configuring Multicast-Only Fast Reroute in a PIM Domain*

List of Sample Output [show multicast route on page 476](#)
[show multicast route \(Bidirectional PIM\) on page 477](#)
[show multicast route brief on page 477](#)
[show multicast route detail on page 478](#)
[show multicast route extensive \(Bidirectional PIM\) on page 478](#)
[show multicast route extensive \(Multicast-Only Fast Reroute\) on page 479](#)
[show multicast route instance <instance-name> on page 479](#)
[show multicast route summary on page 480](#)

Output Fields [Table 54 on page 475](#) describes the output fields for the **show multicast route** command. Output fields are listed in the approximate order in which they appear.

Table 54: show multicast route Output Fields

Field Name	Field Description	Level of Output
family	IPv4 address family (INET) or IPv6 address family (INET6).	All levels
Group	Group address. For any-source multicast routes, for example for bidirectional PIM, the group address includes the prefix length.	All levels
Source	Prefix and length of the source as it is in the multicast forwarding table.	All levels
Incoming interface list	List of interfaces that accept incoming traffic. Only shown for routes that do not use strict RPF-based forwarding, for example for bidirectional PIM.	All levels
Upstream interface	Name of the interface on which the packet with this source prefix is expected to arrive.	All levels
Upstream rpf interface list	When multicast-only fast reroute (MoFRR) is enabled, a PIM router propagates join messages on two upstream RPF interfaces to receive multicast traffic on both links for the same join request.	All levels
Downstream interface list	List of interface names to which the packet with this source prefix is forwarded.	All levels
Number of outgoing interfaces	Total number of outgoing interfaces for each (S,G) entry.	extensive
Session description	Name of the multicast session.	detail extensive

Table 54: show multicast route Output Fields (*continued*)

Field Name	Field Description	Level of Output
Statistics	Rate at which packets are being forwarded for this source and group entry (in Kbps and pps), and number of packets that have been forwarded to this prefix. If one or more of the kilobits per second packet forwarding statistic queries fails or times out, the statistics field displays Forwarding statistics are not available . NOTE: On QFX Series switches, this field does not report valid statistics.	detail extensive
Next-hop ID	Next-hop identifier of the prefix. The identifier is returned by the routing device's Packet Forwarding Engine and is also displayed in the output of the show multicast nexthops command.	detail extensive
Incoming interface list ID	For bidirectional PIM, incoming interface list identifier. Identifiers for interfaces that accept incoming traffic. Only shown for routes that do not use strict RPF-based forwarding, for example for bidirectional PIM.	detail extensive
Upstream protocol	The protocol that maintains the active multicast forwarding route for this group or source. When the show multicast route extensive command is used with the display-origin-protocol option, the field name is only Protocol and not Upstream Protocol . However, this field also displays the protocol that installed the active route.	detail extensive
Route type	Type of multicast route. Values can be (S,G) or (*,G).	summary
Route state	Whether the group is Active or Inactive .	summary extensive
Route count	Number of multicast routes.	summary
Forwarding state	Whether the prefix is pruned or forwarding.	extensive
Cache lifetime/timeout	Number of seconds until the prefix is removed from the multicast forwarding table. A value of never indicates a permanent forwarding entry. A value of forever indicates routes that do not have keepalive times.	extensive
Wrong incoming interface notifications	Number of times that the upstream interface was not available.	extensive
Uptime	Time since the creation of a multicast route.	extensive

Sample Output

show multicast route

```

user@host> show multicast route
Family: INET

Group: 228.0.0.0
Source: 10.255.14.144/32

```

```

Upstream interface: local
Downstream interface list:
  so-1/0/0.0

Group: 239.1.1.1
Source: 10.255.14.144/32
Upstream interface: local
Downstream interface list:
  so-1/0/0.0

Group: 239.1.1.1
Source: 10.255.70.15/32
Upstream interface: so-1/0/0.0
Downstream interface list:
  mt-1/1/0.1081344

Family: INET6

```

show multicast route (Bidirectional PIM)

```

user@host> show multicast route
Family: INET

Group: 224.1.1.0/24
Source: *
Incoming interface list:
  lo0.0 ge-0/0/1.0
Downstream interface list:
  ge-0/0/1.0

Group: 224.1.3.0/24
Source: *
Incoming interface list:
  lo0.0 ge-0/0/1.0 xe-4/1/0.0
Downstream interface list:
  ge-0/0/1.0

Group: 225.1.1.0/24
Source: *
Incoming interface list:
  lo0.0 ge-0/0/1.0
Downstream interface list:
  ge-0/0/1.0

Group: 225.1.3.0/24
Source: *
Incoming interface list:
  lo0.0 ge-0/0/1.0 xe-4/1/0.0
Downstream interface list:
  ge-0/0/1.0

Family: INET6

```

show multicast route brief

The output for the **show multicast route brief** command is identical to that for the **show multicast route** command. For sample output, see [show multicast route on page 476](#) or [show multicast route \(Bidirectional PIM\) on page 477](#).

show multicast route detail

```
user@host> show multicast route detail
Family: INET

Group: 228.0.0.0
  Source: 10.255.14.144/32
  Upstream interface: local
  Downstream interface list:
    so-1/0/0.0
  Session description: Unknown
  Statistics: 8 kbps, 100 pps, 45272 packets
  Next-hop ID: 262142
  Upstream protocol: PIM

Group: 239.1.1.1
  Source: 10.255.14.144/32
  Upstream interface: local
  Downstream interface list:
    so-1/0/0.0
  Session description: Administratively Scoped
  Statistics: 0 kbps, 0 pps, 13404 packets
  Next-hop ID: 262142
  Upstream protocol: PIM

Group: 239.1.1.1
  Source: 10.255.70.15/32
  Upstream interface: so-1/0/0.0
  Downstream interface list:
    mt-1/1/0.1081344
  Session description: Administratively Scoped
  Statistics: 46 kbps, 1000 pps, 921077 packets

  Next-hop ID: 262143
  Upstream protocol: PIM

Family: INET6
```

show multicast route extensive (Bidirectional PIM)

```
user@host> show multicast route extensive
Family: INET

Group: 224.1.1.0/24
  Source: *
  Incoming interface list:
    lo0.0 ge-0/0/1.0
  Downstream interface list:
    ge-0/0/1.0
  Number of outgoing interfaces: 1
  Session description: NOB Cross media facilities
  Statistics: 0 kbps, 0 pps, 0 packets
  Next-hop ID: 2097153
  Incoming interface list ID: 585
  Upstream protocol: PIM
  Route state: Active
  Forwarding state: Forwarding
  Cache lifetime/timeout: forever
  Wrong incoming interface notifications: 0

Group: 224.1.3.0/24
```



```

Source: *
Incoming interface list:
  lo0.0 ge-0/0/1.0 xe-4/1/0.0
Downstream interface list:
  ge-0/0/1.0
Number of outgoing interfaces: 1
Session description: NOB Cross media facilities
Statistics: 0 kbps, 0 pps, 0 packets
Next-hop ID: 2097153
Incoming interface list ID: 589
Upstream protocol: PIM
Route state: Active
Forwarding state: Forwarding
Cache lifetime/timeout: forever
Wrong incoming interface notifications: 0

```

Family: INET6

show multicast route extensive (Multicast-Only Fast Reroute)

```
user@host> show multicast route extensive
```

Instance: master Family: INET

```

Group: 225.1.1.1
Source: 10.0.0.1/32
Upstream rpf interface list:
  fe-1/2/13.0 (P) fe-1/2/14.0 (B)
Downstream interface list:
  fe-1/2/15.0
Session description: Unknown
Forwarding statistics are not available
RPF Next-hop ID: 836
Next-hop ID: 1048585
Upstream protocol: PIM
Route state: Active
Forwarding state: Forwarding
Cache lifetime/timeout: 171 seconds
Wrong incoming interface notifications: 0
Uptime: 00:03:09

```

show multicast route instance <instance-name>

```
user@host> show multicast route instance v1 extensive
```

Instance: v1 Family: INET

```

Group: 224.1.1.1
Source: (null)/0
Upstream interface: fe-1/3/0.111
Downstream interface list:
  lt-0/3/0.42 lt-0/3/0.46 lt-0/3/0.43
Number of outgoing interfaces: 3

```

```

Group: 224.1.1.2
Source: (null)/0
Upstream interface: fe-1/3/0.111
Downstream interface list:
  lt-0/3/0.42 lt-0/3/0.46 lt-0/3/0.43
Number of outgoing interfaces: 3

```

```
Group: 224.1.1.3
```

```
Source: (null)/0
Upstream interface: fe-1/3/0.111
Downstream interface list:
  1t-0/3/0.42 1t-0/3/0.46 1t-0/3/0.43
Number of outgoing interfaces: 3
```

```
Instance: v1 Family: INET6
```

show multicast route summary

```
user@host>show multicast route summary
Instance: master Family: INET
```

Route type	Route state	Route count
(S,G)	Active	2
(S,G)	Inactive	3

```
Instance: master Family: INET6
```

show multicast rpf

Syntax	<pre>show multicast rpf <inet inet6> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <prefix> <summary></pre>
Syntax (EX Series Switch and the QFX Series)	<pre>show multicast rpf <inet inet6> <instance <i>instance-name</i>> <prefix> <summary></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Display information about multicast reverse-path-forwarding (RPF) calculations.
Options	<p>none—Display RPF calculation information for all supported address families.</p> <p>inet inet6—(Optional) Display the RPF calculation information for IPv4 or IPv6 family addresses, respectively.</p> <p>instance <i>instance-name</i>—(Optional) Display information about multicast RPF calculations for a specific multicast instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>prefix—(Optional) Display the RPF calculation information for the specified prefix.</p> <p>summary—(Optional) Display a summary of all multicast RPF information.</p>
Required Privilege Level	view
List of Sample Output	<p>show multicast rpf on page 482</p> <p>show multicast rpf inet6 on page 483</p> <p>show multicast rpf prefix on page 484</p> <p>show multicast rpf summary on page 484</p>

Output Fields Table 55 on page 482 describes the output fields for the **show multicast rpf** command. Output fields are listed in the approximate order in which they appear.

Table 55: show multicast rpf Output Fields

Field Name	Field Description
Instance	Name of the routing instance. (Displayed when multicast is configured within a routing instance.)
Source prefix	Prefix and length of the source as it exists in the multicast forwarding table.
Protocol	How the route was learned.
Interface	Upstream RPF interface. NOTE: The displayed interface information does not apply to bidirectional PIM RP addresses. This is because the show multicast rpf command does not take into account equal-cost paths or the designated forwarder. For accurate upstream RPF interface information, always use the show pim join extensive command when bidirectional PIM is configured.
Neighbor	Upstream RPF neighbor. NOTE: The displayed neighbor information does not apply to bidirectional PIM. This is because the show multicast rpf command does not take into account equal-cost paths or the designated forwarder. For accurate upstream RPF neighbor information, always use the show pim join extensive command when bidirectional PIM is configured.

Sample Output

show multicast rpf

```

user@host> show multicast rpf

Multicast RPF table: inet.0, 12 entries

0.0.0.0/0
  Protocol: Static

10.255.14.132/32
  Protocol: Direct
  Interface: lo0.0

10.255.245.91/32
  Protocol: IS-IS
  Interface: so-1/1/1.0
  Neighbor: 192.168.195.21

127.0.0.1/32
Inactive172.16.0.0/12
Protocol: Static
Interface: fxp0.0

```

```

Neighbor: 192.168.14.254

192.168.0.0/16
Protocol: Static
Interface: fxp0.0
Neighbor: 192.168.14.254

192.168.14.0/24
Protocol: Direct
Interface: fxp0.0

192.168.14.132/32
Protocol: Local

192.168.195.20/30
Protocol: Direct
Interface: so-1/1/1.0

192.168.195.22/32
Protocol: Local

192.168.195.36/30
Protocol: IS-IS
Interface: so-1/1/1.0
Neighbor: 192.168.195.21

```

show multicast rpf inet6

```

user@host> show multicast rpf inet6

Multicast RPF table: inet6.0, 12 entries

::10.255.14.132/128
  Protocol: Direct
  Interface: lo0.0

::10.255.245.91/128
  Protocol: IS-IS
  Interface: so-1/1/1.0
  Neighbor: fe80::2a0:a5ff:fe28:2e8c

::192.168.195.20/126
  Protocol: Direct
  Interface: so-1/1/1.0

::192.168.195.22/128
  Protocol: Local

::192.168.195.36/126
  Protocol: IS-IS
  Interface: so-1/1/1.0
  Neighbor: fe80::2a0:a5ff:fe28:2e8c

::192.168.195.76/126
  Protocol: Direct
  Interface: fe-2/2/0.0

::192.168.195.77/128
  Protocol: Local

```

```
fe80::/64
Protocol: Direct
Interface: so-1/1/1.0

fe80::290:69ff:fe0c:993a/128
Protocol: Local

fe80::2a0:a5ff:fe12:84f/128
Protocol: Direct
Interface: lo0.0

ff02::2/128
Protocol: PIM

ff02::d/128
Protocol: PIM
```

show multicast rpf prefix

```
user@host> show multicast rpf ff02::/16

Multicast RPF table: inet6.0, 13 entries

ff02::2/128
    Protocol: PIM

ff02::d/128
    Protocol: PIM

...
```

show multicast rpf summary

```
user@host> show multicast rpf summary

Multicast RPF table: inet.0, 16 entries
Multicast RPF table: inet6.0, 12 entries
```

show multicast scope

Syntax	show multicast scope <inet inet6> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch and the QFX Series)	show multicast scope <inet inet6> <instance <i>instance-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display administratively scoped IP multicast information.
Options	<p>none—Display standard information about administratively scoped multicast information for all supported address families in all routing instances.</p> <p>inet inet6—(Optional) Display scoped multicast information for IPv4 or IPv6 family addresses, respectively.</p> <p>instance <i>instance-name</i>—(Optional) Display administratively scoped information for a specific multicast instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show multicast scope on page 486 show multicast scope inet on page 486 show multicast scope inet6 on page 486
Output Fields	Table 56 on page 485 describes the output fields for the show multicast scope command. Output fields are listed in the approximate order in which they appear.

Table 56: show multicast scope Output Fields

Field Name	Field Description
Scope name	Name of the multicast scope.
Group Prefix	Range of multicast groups that are scoped.
Interface	Interface that is the boundary of the administrative scope.
Resolve Rejects	Number of kernel resolve rejects.

Sample Output

show multicast scope

```
user@host> show multicast scope
```

Scope name	Group Prefix	Interface	Resolve Rejects
232-net	232.232.0.0/16	fe-0/0/0.1	0
local	239.255.0.0/16	fe-0/0/0.1	0
local	ff05::/16	fe-0/0/0.1	0
larry	ff05::1234/128	fe-0/0/0.1	0

show multicast scope inet

```
user@host> show multicast scope inet
```

Scope name	Group Prefix	Interface	Resolve Rejects
232-net	232.232.0.0/16	fe-0/0/0.1	0
local	239.255.0.0/16	fe-0/0/0.1	0

show multicast scope inet6

```
user@host> show multicast scope inet6
```

Scope name	Group Prefix	Interface	Resolve Rejects
local	ff05::/16	fe-0/0/0.1	0
larry	ff05::1234/128	fe-0/0/0.1	0

show multicast snooping next-hops

Syntax	show multicast snooping next-hops <brief detail>
Release Information	Command introduced in Junos OS Release 11.2.
Description	Display information about the IP multicast snooping next-hops.
Options	brief detail —(Optional) Display the specified level of output.
Required Privilege Level	view
List of Sample Output	show multicast snooping next-hops on page 488
Output Fields	Table 57 on page 487 describes the output fields for the show multicast snooping next-hops command. Output fields are listed in the approximate order in which they appear.

Table 57: show multicast snooping next-hops Output Fields

Field Name	Field Description
Family	Protocol family for which multicast snooping next hops are displayed: INET or INET6 .
Refcount	Number of cache entries that are using this next hop.
KRefcount	Kernel reference count for the next hop.
Downstream interface	Interface names associated with each multicast next-hop ID.
Nexthop Id	Identifier for the next-hop.

Sample Output

show multicast snooping next-hops

```
user@host> show multicast snooping next-hops
Family: INET
ID          Refcount KRefCount Downstream interface Nexthop Id
1048574      4          1 ge-0/1/0.1000
              ge-0/1/2.1000
              ge-0/1/3.1000

1048574      4          1 ge-0/1/0.1000-(2000)
              1048575
              1048576

1048575      2          0 ge-0/1/2.1000-(2001)
              ge-0/1/3.1000-(2002)

1048576      2          0 lsi.1048578-(2003)
              lsi.1048579-(2004)
```

show multicast sessions

Syntax	show multicast sessions <brief detail extensive> <logical-system (all <i>logical-system-name</i>)> < <i>regular-expression</i> >
Syntax (EX Series Switch and the QFX Series)	show multicast sessions <brief detail extensive> < <i>regular-expression</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display information about announced IP multicast sessions.
Options	<p>none—Display standard information about all multicast sessions for all routing instances.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>regular-expression</i>—(Optional) Display information about announced sessions that match a UNIX-style regular expression.</p>
Required Privilege Level	view
List of Sample Output	show multicast sessions on page 490 show multicast sessions <i>regular-expression</i> detail on page 490
Output Fields	<p>Table 58 on page 489 describes the output fields for the show multicast sessions command. Output fields are listed in the approximate order in which they appear.</p>

Table 58: show multicast sessions Output Fields

Field Name	Field Description
<i>session-name</i>	Name of the known announced multicast sessions.

Sample Output

show multicast sessions

```
user@host> show multicast sessions
1-Department of Biological Sciences, LSU
...
Monterey Bay - DockCam
Monterey Bay - JettyCam
Monterey Bay - StandCam
Monterey DockCam
Monterey DockCam / ROV cam
...
NASA TV (MPEG-1)
...
UO Broadcast - NASA Videos - 25 Years of Progress
UO Broadcast - NASA Videos - Journey through the Solar System
UO Broadcast - NASA Videos - Life in the Universe
UO Broadcast - NASA Videos - Nasa and the Airplane
UO Broadcasts OPB's Oregon Story
UO DOD News Clips
UO Medical Management of Biological Casualties (1)
UO Medical Management of Biological Casualties (2)
UO Medical Management of Biological Casualties (3)
...
376 active sessions.
```

show multicast sessions regular-expression detail

```
user@host> show multicast sessions "NASA TV" detail
SDP Version: 0  Originated by: -@128.223.83.33
Session: NASA TV (MPEG-1)
Description: NASA television in MPEG-1 format, provided by Private University.
Please contact the UO if you have problems with this feed.
Email: Your Name Here <multicast@lists.private.edu>
Phone: Your Name Here <888/555-1212>
Bandwidth: AS:1000
Start time: permanent
Stop time: none
Attribute: type:broadcast
Attribute: tool:IP/TV Content Manager 3.4.14
Attribute: live:capture:1
Attribute: x-iptv-capture:mp1s
Media: video 54302 RTP/AVP 32 31 96 97
Connection Data: 224.2.231.45 ttl 127
Attribute: quality:8
Attribute: framerate:30
Attribute: rtpmap:96 WBIH/90000
Attribute: rtpmap:97 MP4V-ES/90000
Attribute: x-iptv-svr:video 128.223.91.191 live
Attribute: fmtp:32 type=mpeg1
Media: audio 28848 RTP/AVP 14 0 96 3 5 97 98 99 100 101 102 10 11 103 104 105 106
Connection Data: 224.2.145.37 ttl 127
Attribute: rtpmap:96 X-WAVE/8000
Attribute: rtpmap:97 L8/8000/2
Attribute: rtpmap:98 L8/8000
Attribute: rtpmap:99 L8/22050/2
Attribute: rtpmap:100 L8/22050
Attribute: rtpmap:101 L8/11025/2
Attribute: rtpmap:102 L8/11025
Attribute: rtpmap:103 L16/22050/2
```

Attribute: rtpmap:104 L16/22050

1 matching sessions.

show multicast snooping route

Syntax show multicast snooping route
 <brief | detail | extensive>
 <active | all | inactive>
 <bridge-domain *bridge-domain-name*>
 <control>
 <data>
 <group *group*>
 <instance *instance-name*>
 <mesh-group *mesh-group-name*>
 <qualified-vlan *vlan-id*>
 <*regular-expression*>
 <source-prefix *source-prefix*>
 <vlan *vlan-id*>

Release Information Command introduced in Junos OS Release 8.5.
 Support for **control**, **data**, **qualified-vlan** and **vlan** options introduced in Junos OS Release 13.3 for EX Series switches.

Description Display the entries in the IP multicast snooping forwarding table. You can display some of this information with the **show route table inet.1** command.

Options **none**—Display standard information about all entries in the multicast snooping table for all virtual switches and all bridge domains.

brief | detail | extensive—(Optional) Display the specified level of output.

active | all | inactive—(Optional) Display all active entries, all entries, or all inactive entries, respectively, in the multicast snooping table.

bridge-domain *bridge-domain*—(Optional) Display the entries for a particular bridge domain.

control—(Optional) Display control route entries.

data—(Optional) Display data route entries.

group *group*—(Optional) Display the entries for a particular group.

instance *instance-name*—(Optional) Display the entries for a multicast instance.

mesh-group *mesh-group-name*—(Optional) Display the entries for a particular mesh group.

qualified-vlan *vlan-id*—(Optional) Display the entries for a particular qualified VLAN.

regular-expression—(Optional) Display information about the multicast forwarding table entries that match a UNIX-style regular expression.

source-prefix *source-prefix*—(Optional) Display the entries for a particular source prefix.

vlan *vlan-id*—(Optional) Display the entries for a particular VLAN.

Required Privilege Level view

List of Sample Output [show multicast snooping route bridge-domain on page 493](#)
[show multicast snooping route instance vs on page 494](#)

Output Fields [Table 59 on page 493](#) describes the output fields for the **show multicast snooping route** command. Output fields are listed in the approximate order in which they appear.

Table 59: show multicast snooping route Output Fields

Field Name	Field Description	Level of Output
Nexthop Bulking	Displays whether next-hop bulk updating is ON or OFF (only for routing-instance type of virtual switch or vpls).	All levels
Family	IPv4 address family (INET) or IPv6 address family (INET6).	All levels
Group	Group address.	All levels
Source	Prefix and length of the source as it is in the multicast forwarding table.	All levels
Routing-instance	Name of the routing instance to which this routing information applies. (Displayed when multicast is configured within a routing instance.)	All levels
Learning Domain	Name of the learning domain to which this routing information applies.	detail extensive
Statistics	Rate at which packets are being forwarded for this source and group entry (in Kbps and pps), and number of packets that have been forwarded to this prefix.	detail extensive
Next-hop ID	Next-hop identifier of the prefix. The identifier is returned by the router's Packet Forwarding Engine and is also displayed in the output of the show multicast nexthops command.	detail extensive
Route state	Whether the group is Active or Inactive .	extensive
Forwarding state	Whether the prefix is Pruned or Forwarding .	extensive
Cache lifetime/timeout	Number of seconds until the prefix is removed from the multicast forwarding table. A value of never indicates a permanent forwarding entry.	extensive

Sample Output

[show multicast snooping route bridge-domain](#)

```

user@host> show multicast snooping route bridge-domain br-dom-1 extensive
Family: INET

Group: 232.1.1.1
Source: 192.168.3.100/32
Downstream interface list:
    ge-0/1/0.200
Statistics: 0 kbps, 0 pps, 1 packets
Next-hop ID: 1048577

```

Route state: Active
Forwarding state: Forwarding
Cache lifetime/timeout: 240 seconds

show multicast snooping route instance vs

```
user@host> show multicast snooping route instance vs
NextHop Bulking: ON
```

Family: INET

Group: 224.0.0.0
Bridge-domain: vsid500

Group: 225.1.0.1
Bridge-domain: vsid500
Downstream interface list: vsid500
ge-0/3/8.500 ge-1/1/9.500 ge1/2/5.500

show multicast statistics

Syntax	show multicast statistics <inet inet6> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display IP multicast statistics.
Options	<p>none—Display multicast statistics for all supported address families for all routing instances.</p> <p>inet inet6—(Optional) Display multicast statistics for IPv4 or IPv6 family addresses, respectively.</p> <p>instance <i>instance-name</i>—(Optional) Display statistics for a specific routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Additional Information	The input and output interface multicast statistics are consistent, but not timely. They are constructed from the forwarding statistics, which are gathered at 30-second intervals. Therefore, the output from this command always lags the true count by up to 30 seconds.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear multicast statistics on page 381
List of Sample Output	show multicast statistics on page 497
Output Fields	Table 60 on page 495 describes the output fields for the show multicast statistics command. Output fields are listed in the approximate order in which they appear.

Table 60: show multicast statistics Output Fields

Field Name	Field Description
Family	Protocol family for which multicast statistics are displayed: INET or INET6 .
Interface	Name of the interface for which statistics are being reported.
Routing Protocol	Primary multicast protocol on the interface: PIM , DVMRP for INET , or PIM for INET6 .
Mismatch	Number of multicast packets that did not arrive on the correct upstream interface.
Kernel Resolve	Number of resolve requests processed by the primary multicast protocol on the interface.

Table 60: show multicast statistics Output Fields (*continued*)

Field Name	Field Description
Resolve No Route	Number of resolve requests that were ignored because there was no route to the source.
In Kbytes	Total accumulated incoming packets (in KB) since the last time the clear multicast statistics command was issued.
Out Kbytes	Total accumulated outgoing packets (in KB) since the last time the clear multicast statistics command was issued.
Mismatch error	Number of mismatches that were ignored because of internal errors.
Mismatch No Route	Number of mismatches that were ignored because there was no route to the source.
Routing Notify	Number of times that the multicast routing system has been notified of a new multicast source by a multicast routing protocol .
Resolve Error	Number of resolve requests that were ignored because of internal errors.
In Packets	Total number of incoming packets since the last time the clear multicast statistics command was issued.
Out Packets	Total number of outgoing packets since the last time the clear multicast statistics command was issued.
Resolve requests on interfaces not enabled for multicast <i>n</i>	Number of resolve requests on interfaces that are not enabled for multicast that have accumulated since the clear multicast statistics command was last issued.
Resolve requests with no route to source <i>n</i>	Number of resolve requests with no route to the source that have accumulated since the clear multicast statistics command was last issued.
Routing notifications on interfaces not enabled for multicast <i>n</i>	Number of routing notifications on interfaces not enabled for multicast that have accumulated since the clear multicast statistics command was last issued.
Routing notifications with no route to source <i>n</i>	Number of routing notifications with no route to the source that have accumulated since the clear multicast statistics command was last issued.
Interface Mismatches on interfaces not enabled for multicast <i>n</i>	Number of interface mismatches on interfaces not enabled for multicast that have accumulated since the clear multicast statistics command was last issued.
Group Membership on interfaces not enabled for multicast <i>n</i>	Number of group memberships on interfaces not enabled for multicast that have accumulated since the clear multicast statistics command was last issued.

Sample Output

show multicast statistics

```

user@host> show multicast statistics
Address family: INET
Interface: fe-0/0/0
  Routing Protocol:      PIM  Mismatch error:      0
  Mismatch:              0    Mismatch No Route:    0
  Kernel Resolve:        10    Routing Notify:       0
  Resolve No Route:      0    Resolve Error:        0
  In Kbytes:             4641  In Packets:           50454
  Out Kbytes:            0    Out Packets:          0
Interface: so-0/1/1.0
  Routing Protocol:      PIM  Mismatch error:      0
  Mismatch:              0    Mismatch No Route:    0
  Kernel Resolve:        0    Routing Notify:       0
  Resolve No Route:      0    Resolve Error:        0
  In Kbytes:             0    In Packets:           0
  Out Kbytes:            4641  Out Packets:          50454

Resolve requests on interfaces not enabled for multicast 0
Resolve requests with no route to source 0
Routing notifications on interfaces not enabled for multicast 0
Routing notifications with no route to source 0
Interface Mismatches on interfaces not enabled for multicast 0
Group Membership on interfaces not enabled for multicast 25

Address family: INET6
Interface: fe-0/0/0.0
  Routing Protocol:      PIM  Mismatch error:      0
  Mismatch:              0    Mismatch No Route:    0
  Kernel Resolve:        0    Routing Notify:       0
  Resolve No Route:      0    Resolve Error:        0
  In Kbytes:             0    In Packets:           0
  Out Kbytes:            0    Out Packets:          0
Interface: so-0/1/1.0
  Routing Protocol:      PIM  Mismatch error:      0
  Mismatch:              0    Mismatch No Route:    0
  Kernel Resolve:        0    Routing Notify:       0
  Resolve No Route:      0    Resolve Error:        0
  In Kbytes:             0    In Packets:           0
  Out Kbytes:            0    Out Packets:          0

Resolve requests on interfaces not enabled for multicast 0
Resolve requests with no route to source 0
Routing notifications on interfaces not enabled for multicast 0
Routing notifications with no route to source 0
Interface Mismatches on interfaces not enabled for multicast 0
Group Membership on interfaces not enabled for multicast 0

```

show multicast usage

Syntax	<pre>show multicast usage <brief detail> <inet inet6> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switch and the QFX Series)	<pre>show multicast usage <brief detail> <inet inet6> <instance <i>instance-name</i>></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Display usage information about the 10 most active Distance Vector Multicast Routing Protocol (DVMRP) or Protocol Independent Multicast (PIM) groups.
Options	<p>none—Display multicast usage information for all supported address families for all routing instances.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>inet inet6—(Optional) Display usage information for IPv4 or IPv6 family addresses, respectively.</p> <p>instance <i>instance-name</i>—(Optional) Display information about the most active DVMRP or PIM groups for a specific multicast instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	<p>show multicast usage on page 499</p> <p>show multicast usage brief on page 499</p> <p>show multicast usage instance on page 499</p> <p>show multicast usage detail on page 500</p>
Output Fields	<p>Table 61 on page 498 describes the output fields for the show multicast usage command. Output fields are listed in the approximate order in which they appear.</p>

Table 61: show multicast usage Output Fields

Field Name	Field Description
Instance	Name of the routing instance. (Displayed when multicast is configured within a routing instance.)

Table 61: show multicast usage Output Fields (*continued*)

Field Name	Field Description
Group	Group address.
Sources	Number of sources.
Packets	Number of packets that have been forwarded to this prefix. If one or more of the packets forwarded statistic queries fails or times out, the packets field displays unavailable .
Bytes	Number of bytes that have been forwarded to this prefix. If one or more of the packets forwarded statistic queries fails or times out, the bytes field displays unavailable .
Prefix	IP address.
/len	Prefix length.
Groups	Number of multicast groups.

Sample Output

show multicast usage

```

user@host> show multicast usage
Group          Sources  Packets      Bytes
228.0.0.0      1        52847      4439148
239.1.1.1      2        13450      1125530

Prefix         /len  Groups  Packets      Bytes
10.255.14.144  /32   2        66254      5561304
10.255.70.15   /32   1         43        3374...
```

show multicast usage brief

The output for the **show multicast usage brief** command is identical to that for the **show multicast usage** command. For sample output, see [show multicast usage on page 499](#).

show multicast usage instance

```

user@host> show multicast usage instance VPN-A
Group          Sources  Packets      Bytes
224.2.127.254  1        5538      509496
224.0.1.39     1         13        624
224.0.1.40     1         13        624

Prefix         /len  Groups  Packets      Bytes
192.168.195.34 /32   1        5538      509496
10.255.14.30   /32   1         13        624
10.255.245.91  /32   1         13        624
...
```

show multicast usage detail

```
user@host> show multicast usage detail
```

Group	Sources	Packets	Bytes
228.0.0.0	1	53159	4465356
Source: 10.255.14.144 /32 Packets: 53159 Bytes: 4465356			
239.1.1.1	2	13450	1125530
Source: 10.255.14.144 /32 Packets: 13407 Bytes: 1122156			
Source: 10.255.70.15 /32 Packets: 43 Bytes: 3374			

Prefix	/len	Groups	Packets	Bytes
10.255.14.144	/32	2	66566	5587512
Group: 228.0.0.0		Packets: 53159	Bytes: 4465356	
Group: 239.1.1.1		Packets: 13407	Bytes: 1122156	
10.255.70.15	/32	1	43	3374
Group: 239.1.1.1		Packets: 43	Bytes: 3374	

show pgm negative-acknowledgments

Syntax	show pgm negative-acknowledgments
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display the sent or received Pragmatic General Multicast (PGM) negative acknowledgments (NAKs), the source-path message (SPM) sequence number being negatively acknowledged, and the current state of repair.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show pgm negative-acknowledgments on page 502
Output Fields	Table 62 on page 501 describes the output fields for the show pgm negative-acknowledgments command. Output fields are listed in the approximate order in which they appear.

Table 62: show pgm negative-acknowledgments Output Fields

Field Name	Field Description
Global source id	Global source identifier (GSI), which combines with the source port to determine the transport session identifier (TSI).
Network layer address	Network layer address of the local system.
Source port	Source port number, which is combined with the GSI to determine the TSI.
SPM sequence number	Numeric sequence identifier of the source-path message.
Window (trailing/leading sequence)	Range of sequence numbers used by the source for sequentially numbering and transmitting the most recent packets. The trailing (or left) edge of the transmit window is the sequence number of the oldest data packet available for repair from a source. The leading (or right) edge of the transmit window is defined as the sequence number of the most recent data packet a source has transmitted.
Outstanding NAKS	<p>Total number of outstanding negative acknowledgments sent or received by the local system. NAK packets indicate that a packet in the expected original data sequence has been detected as missing.</p> <ul style="list-style-type: none"> • Sequence number—Numeric sequence identifier of the source-path message. • Group—Group address. • Source—Multicast source. • Interface—Interface name. • Receiver—IP address receiving the multicast.

Sample Output

show pgm negative-
acknowledgments

```
user@host> show pgm negative-acknowledgments
Global source ID: 010203040506 Source port: 1111
  Network layer address: 10.38.0.1
  SPM sequence number: 1
  Window (trailing/leading sequence): 0/1
  Outstanding NAKs:
    Sequence number: 1
    Group: 225.1.1.1
    Source: 192.168.195.121
    Interface: t3-0/2/0:0 Receiver: 10.38.0.10
```


show pgm source-path-messages

Syntax	show pgm source-path-messages
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display the Pragmatic General Multicast (PGM) source-path messages received.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show pgm source-path-messages on page 503
Output Fields	Table 63 on page 503 describes the output fields for the show pgm source-path-messages command. Output fields are listed in the approximate order in which they appear.

Table 63: show pgm source-path-messages Output Fields

Field Name	Field Description
Global source ID	Global source identifier (GSI), which combines with the source port to determine the transport session identifier (TSI).
Port	Source port number, which combines with the GSI to determine the TSI.
SPM number	Numeric sequence identifier of the source-path message.
Trail number	Sequence number of the oldest data packet available for repair from a source.
Lead number	Sequence number of the most recent data packet a source has transmitted.
Network layer address	Network layer address of the local system.

Sample Output

show pgm source-path-messages

```

user@host> show pgm source-path-messages
Global source ID  Port  SPM number  Trail number  Lead number  Network layer address
010203040506     1111         1           0             1    10.38.0.1

```

show pgm statistics

Syntax	show pgm statistics
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display Pragmatic General Multicast (PGM) packet statistics, including general loss and repair statistics.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show pgm statistics on page 506
Output Fields	Table 64 on page 504 describes the output fields for the show pgm statistics command. Output fields are listed in the approximate order in which they appear.

Table 64: show pgm statistics Output Fields

Field Name	Field Description
PGM type, # received, # sent	<p>Number of packets received and sent for the following PGM packet types:</p> <ul style="list-style-type: none"> SPM—Number of total source path messages received and sent by the local system. Source path messages (SPMs) are sent by a source to establish the source path state in network elements and to provide the transmit-window state to receivers. POLL—Total number of poll requests received and sent by the local system. POLR—Total number of poll responses received and sent by the local system. ODATA—Total number of original data packets received and sent by the local system. RDATA—Total number of repair data packets received and sent by the local system. RDATA packets are generated in response to negative acknowledgments (NAKs), which indicate a missing packet from the original data sequence. NAK—Total number of negative acknowledgments received and sent by the local system. NAK packets indicate that a packet in the expected original data sequence has been detected as missing. NULLNAK—Total number of null negative acknowledgments received and sent by the local system. NULLNAKs are transmitted by a designated local repairer that receives NAKs redirected to it by either receivers or network elements to provide flow-control feedback to a source. NCF—Total number of NAK confirmations received and sent by the local system. NAK confirmations are generated in response to NAK packets that are received. SPMR—Total number of source path message requests (SPMRs) received and sent by the local system. SPMRs are used to solicit a source path message from a source in a nonimplosive way. The typical application is for late-joining receivers to solicit source path messages directly from a source in order to be able to send NAKs for missing packets, without having to wait for a regularly scheduled source path message from that source. OTHER—Total number of other PGM packets received and sent by the local system.
packets shorter than minimum PGM header length	Total number of packets received with headers that are shorter than the minimum required PGM header length.

Table 64: show pgm statistics Output Fields (*continued*)

Field Name	Field Description
packets received with incorrect check sum	Total number of packets received with an incorrect checksum. The checksum field is the 1's complement of the 1's complement sum of the entire PGM packet, including the header.
packets received with zero check sum	Total number of packets received with a zero checksum. If the computed checksum is zero, it is transmitted as all ones. A value of zero in this field means that the transmitter generated no checksum.
packets received with TSDU length incorrect	Total number of packets received with an incorrect Transport Service Data Unit (TSDU) length (16 bits).
packets received with SPM length incorrect	Total number of packets received with an incorrect source path message length.
packets received with unknown SPM address family	Total number of packets received with an unknown source path message address family indicator (AFI).
packets received with NAK length incorrect	Total number of packets received with an incorrect NAK length.
packets received with unknown NAK address family	Total number of packets received with an unknown NAK address family indicator (AFI).
packets received with NAK for unknown TSI	Total number of NAK packets received with an unknown transport session identifier (TSI).
packets received when NAK throttled	Total number of packets received when NAK is throttled.
packets received with NCF length incorrect	Total number of packets received with an incorrect NAK confirmation length.
packets received with unknown NCF address family	Total number of packets received with an unknown NAK confirmation address family indicator (AFI).
packets received with NCF for unknown TSI	Total number of NAK confirmation packets received with an unknown transport session identifier (TSI).
packets received with RDATA length incorrect	Total number of packets received with an incorrect RDATA length.
packets received with RDATA for unknown TSI	Total number of RDATA packets received with an unknown transport session identifier (TSI).

Sample Output

show pgm statistics

```
user@host> show pgm statistics
PGM type      # received    # sent
SPM            0             0
POLL           0             0
POLR           0             0
ODATA          0             0
RDATA          0             0
NAK            0             0
NULLNAK        0             0
NCF            0             0
SPMR           0             0
OTHER          0             0

packets shorter than minimum PGM header length :      0
packets received with incorrect check sum       :      0
packets received with zero check sum            :      0
packets received with TSDU length incorrect     :      0
packets received with SPM length incorrect      :      0
packets received with unknown SPM address family:      0
packets received with NAK length incorrect      :      0
packets received with unknown NAK address family:      0
packets received with NAK for unknown TSI       :      0
packets received when NAK throttled             :      0
packets received with NCF length incorrect      :      0
packets received with unknown NCF address family:      0
packets received with NCF for unknown TSI       :      0
packets received with RDATA length incorrect    :      0
packets received with RDATA for unknown TSI     :      0
```

show pim bootstrap

Syntax	show pim bootstrap <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch and the QFX Series)	show pim bootstrap <instance <i>instance-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. instance option introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	For sparse mode only, display information about Protocol Independent Multicast (PIM) bootstrap routers.
Options	<p>none—Display PIM bootstrap router information for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Display information about bootstrap routers for a specific PIM-enabled routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show pim bootstrap on page 508 show pim bootstrap instance on page 508
Output Fields	Table 65 on page 507 describes the output fields for the show pim bootstrap command. Output fields are listed in the approximate order in which they appear.

Table 65: show pim bootstrap Output Fields

Field Name	Field Description
Instance	Name of the routing instance.
BSR	Bootstrap router.
Pri	Priority of the routing device as elected to be the bootstrap router.
Local address	Local routing device address.
Pri	Local routing device address priority to be elected as the bootstrap router.
State	Local routing device election state: Candidate , Elected , or Ineligible .

Table 65: show pim bootstrap Output Fields (*continued*)

Field Name	Field Description
Timeout	How long until the local routing device declares the bootstrap router to be unreachable, in seconds.

Sample Output

show pim bootstrap

```
user@host> show pim bootstrap
Instance: PIM.master
```

BSR	Pri	Local address	Pri	State	Timeout
None	0	10.255.71.46	0	InEligible	0
feco:1:1:1:1:0:aff:785c	34	feco:1:1:1:1:0:aff:7c12	0	InEligible	0

show pim bootstrap instance

```
user@host> show pim bootstrap instance VPN-A
Instance: PIM.VPN-A
```

BSR	Pri	Local address	Pri	State	Timeout
None	0	192.168.196.105	0	InEligible	0

show pim interfaces

Syntax	show pim interfaces <inet inet6> <instance (<i>instance-name</i> all)> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch and the QFX Series)	show pim interfaces <inet inet6> <instance (<i>instance-name</i> all)>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. Support for bidirectional PIM added in Junos OS Release 12.1. Support for the instance all option added in Junos OS Release 12.1.
Description	Display information about the interfaces on which Protocol Independent Multicast (PIM) is configured.
Options	<p>none—Display interface information for all family addresses for the main instance.</p> <p>inet inet6—(Optional) Display interface information for IPv4 or IPv6 family addresses, respectively.</p> <p>instance (<i>instance-name</i> all)—(Optional) Display information about interfaces for a specific PIM-enabled routing instance or for all routing instances.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show pim interfaces on page 510
Output Fields	Table 66 on page 509 describes the output fields for the show pim interfaces command. Output fields are listed in the approximate order in which they appear.

Table 66: show pim interfaces Output Fields

Field Name	Field Description
Instance	Name of the routing instance.
Name	Interface name.
State	State of the interface. The state also is displayed in the show interfaces command.

Table 66: show pim interfaces Output Fields (*continued*)

Field Name	Field Description
Mode	<p>PIM mode running on the interface:</p> <ul style="list-style-type: none"> • B—In bidirectional mode, multicast groups are carried across the network over bidirectional shared trees. This type of tree minimizes PIM routing state, which is especially important in networks with numerous and dispersed senders and receivers. • S—In sparse mode, routing devices must join and leave multicast groups explicitly. Upstream routing devices do not forward multicast traffic to this routing device unless this device has sent an explicit request (using a join message) to receive multicast traffic. • Dense—Unlike sparse mode, where data is forwarded only to routing devices sending an explicit request, dense mode implements a flood-and-prune mechanism, similar to DVMRP (the first multicast protocol used to support the multicast backbone). (Not supported on QFX Series.) • Sparse-Dense—Sparse-dense mode allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as dense is not mapped to a rendezvous point (RP). Instead, data packets destined for that group are forwarded using PIM-Dense Mode (PIM-DM) rules. A group specified as sparse is mapped to an RP, and data packets are forwarded using PIM-Sparse Mode (PIM-SM) rules. (Not supported on QFX Series.) <p>When sparse-dense mode is configured, the output includes both S and D. When bidirectional-sparse mode is configured, the output includes S and B. When bidirectional-sparse-dense mode is configured, the output includes B, S, and D.</p>
IP	Version number of the address family on the interface: 4 (IPv4) or 6 (IPv6).
V	PIM version running on the interface: 1 or 2.
State	<p>State of PIM on the interface:</p> <ul style="list-style-type: none"> • Active—Bidirectional mode is enabled on the interface and on all PIM neighbors. • DR—Designated router. • NotCap—Bidirectional mode is not enabled on the interface. This can happen when bidirectional PIM is not configured locally, when one of the neighbors is not configured for bidirectional PIM, or when one of the neighbors has not implemented the bidirectional PIM protocol. • NotDR—Not the designated router. • P2P—Point to point.
NbrCnt	Number of neighbors that have been seen on the interface.
JoinCnt(sg)	Number of (s,g) join messages that have been seen on the interface.
JointCnt(*g)	Number of (*g) join messages that have been seen on the interface.
DR address	Address of the designated router.

Sample Output

show pim interfaces

```

user@host> show pim interfaces
Stat = Status, V = Version, NbrCnt = Neighbor Count,
S = Sparse, D = Dense, B = Bidirectional,
DR = Designated Router, P2P = Point-to-point link,

```


Active = Bidirectional is active, NotCap = Not Bidirectional Capable

Name	Stat	Mode	IP	V	State	NbrCnt	JoinCnt(sg/*g)	DR address
ge-0/3/0.0	Up	S	4	2	NotDR,NotCap	1	0/0	40.0.0.3
ge-0/3/3.50	Up	S	4	2	DR,NotCap	1	9901/100	50.0.0.2
ge-0/3/3.51	Up	S	4	2	DR,NotCap	1	0/0	51.0.0.2
pe-1/2/0.32769	Up	S	4	2	P2P,NotCap	0	0/0	

show pim join

Syntax	<pre>show pim join <brief detail extensive summary> <bidirectional dense sparse> <exact> <inet inet6> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <range> <rp <i>ip-address/prefix</i> source <i>ip-address/prefix</i>> <sg star-g></pre>
Syntax (EX Series Switch and the QFX Series)	<pre>show pim join <brief detail extensive summary> <dense sparse> <exact> <inet inet6> <instance <i>instance-name</i>> <range> <rp <i>ip-address/prefix</i> source <i>ip-address/prefix</i>> <sg star-g></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>summary option introduced in Junos OS Release 9.6.</p> <p>inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Multiple new filter options introduced in Junos OS Release 13.2.</p>
Description	<p>Display information about Protocol Independent Multicast (PIM) groups for all PIM modes.</p> <p>For bidirectional PIM, display information about PIM group ranges (*G-range) for each active bidirectional RP group range, in addition to each of the joined (*G) routes.</p>
Options	<p>none—Display the standard information about PIM groups for all supported family addresses for all routing instances.</p> <p>brief detail extensive summary—(Optional) Display the specified level of output.</p> <p>bidirectional dense sparse—(Optional) Display information about PIM bidirectional mode, dense mode, or sparse and source-specific multicast (SSM) mode entries.</p> <p>exact—(Optional) Display information about only the group that exactly matches the specified group address.</p> <p>inet inet6—(Optional) Display PIM group information for IPv4 or IPv6 family addresses, respectively.</p>

instance *instance-name*—(Optional) Display information about groups for the specified PIM-enabled routing instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

range—(Optional) Address range of the group, specified as *prefix/prefix-length*.

rp *ip-address/prefix* | source *ip-address/prefix*—(Optional) Display information about the PIM entries with a specified rendezvous point (RP) address and prefix or with a specified source address and prefix. You can omit the prefix.

sg | star-g—(Optional) Display information about PIM (S,G) or (*G) entries.

Required Privilege Level

view

Related Documentation

- [clear pim join on page 385](#)
- *Example: Configuring Multicast-Only Fast Reroute in a PIM Domain*
- *Example: Configuring Bidirectional PIM*
- *Example: Configuring PIM State Limits*

List of Sample Output

[show pim join summary on page 517](#)
[show pim join \(PIM Sparse Mode\) on page 517](#)
[show pim join \(Bidirectional PIM\) on page 518](#)
[show pim join inet6 on page 518](#)
[show pim join inet6 star-g on page 519](#)
[show pim join instance <instance-name> on page 519](#)
[show pim join detail on page 519](#)
[show pim join extensive \(PIM Sparse Mode\) on page 520](#)
[show pim join extensive \(Bidirectional PIM\) on page 521](#)
[show pim join extensive \(Bidirectional PIM with a Directly Connected Phantom RP\) on page 522](#)
[show pim join instance <instance-name> extensive on page 522](#)
[show pim join extensive \(Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs\) on page 523](#)
[show pim join extensive \(Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs\) on page 524](#)
[show pim join summary on page 525](#)
[show pim join \(PIM Sparse Mode\) on page 526](#)
[show pim join \(Bidirectional PIM\) on page 526](#)
[show pim join inet6 on page 527](#)
[show pim join inet6 star-g on page 527](#)
[show pim join instance <instance-name> on page 527](#)
[show pim join detail on page 528](#)
[show pim join extensive \(PIM Sparse Mode\) on page 528](#)
[show pim join extensive \(Bidirectional PIM\) on page 529](#)

[show pim join extensive \(Bidirectional PIM with a Directly Connected Phantom RP\) on page 530](#)

[show pim join instance <instance-name> extensive on page 530](#)

[show pim join extensive \(Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs\) on page 531](#)

[show pim join extensive \(Multipoint LDP with Multicast-Only Fast Reroute\) on page 532](#)

Output Fields [Table 67 on page 514](#) describes the output fields for the **show pim join** command. Output fields are listed in the approximate order in which they appear.

Table 67: show pim join Output Fields

Field Name	Field Description	Level of Output
Instance	Name of the routing instance.	brief detail extensive summary none
Family	Name of the address family: inet (IPv4) or inet6 (IPv6).	brief detail extensive summary none
Route type	Type of multicast route: (S,G) or (*,G).	summary
Route count	Number of (S,G) routes and number of (*,G) routes.	summary
R	Rendezvous Point Tree.	brief detail extensive none
S	Sparse.	brief detail extensive none
W	Wildcard.	brief detail extensive none
Group	Group address.	brief detail extensive none
Bidirectional group prefix length	For bidirectional PIM, length of the IP prefix for RP group ranges.	All levels
Source	Multicast source: <ul style="list-style-type: none"> • * (wildcard value) • <i>ipv4-address</i> • <i>ipv6-address</i> 	brief detail extensive none
RP	Rendezvous point for the PIM group.	brief detail extensive none
Flags	PIM flags: <ul style="list-style-type: none"> • bidirectional—Bidirectional mode entry. • dense—Dense mode entry. • rptree—Entry is on the rendezvous point tree. • sparse—Sparse mode entry. • spt—Entry is on the shortest-path tree for the source. • wildcard—Entry is on the shared tree. 	brief detail extensive none

Table 67: show pim join Output Fields (*continued*)

Field Name	Field Description	Level of Output
Upstream interface	<p>RPF interface toward the source address for the source-specific state (S,G) or toward the rendezvous point (RP) address for the non-source-specific state (*G).</p> <p>For bidirectional PIM, RP Link means that the interface is directly connected to a subnet that contains a phantom RP address.</p> <p>A pseudo multipoint LDP (M-LDP) interface appears on egress nodes in M-LDP point-to-multipoint LSPs with inband signaling.</p>	brief detail extensive none
Upstream neighbor	<p>Information about the upstream neighbor: Direct, Local, Unknown, or a specific IP address.</p> <p>For bidirectional PIM, Direct means that the interface is directly connected to a subnet that contains a phantom RP address.</p> <p>The multipoint LDP (M-LDP) root appears on egress nodes in M-LDP point-to-multipoint LSPs with inband signaling.</p>	extensive
Active upstream interface	When multicast-only fast reroute (MoFRR) is configured in a PIM domain, the upstream interface for the active path. A PIM router propagates join messages on two upstream RPF interfaces to receive multicast traffic on both links for the same join request. Preference is given to two paths that do not converge to the same immediate upstream router. PIM installs appropriate multicast routes with upstream neighbors as RPF next hops with two (primary and backup) interfaces.	extensive
Active upstream neighbor	On the MoFRR primary path, the IP address of the neighbor that is directly connected to the active upstream interface.	extensive
MoFRR Backup upstream interface	<p>The MoFRR upstream interface that is used when the primary path fails.</p> <p>When the primary path fails, the backup path is upgraded to primary, and traffic is forwarded accordingly. If there are alternate paths available, a new backup path is calculated and the appropriate multicast route is updated or installed.</p>	extensive

Table 67: show pim join Output Fields (*continued*)

Field Name	Field Description	Level of Output
Upstream state	<p>Information about the upstream interface:</p> <ul style="list-style-type: none"> • Join to RP—Sending a join to the rendezvous point. • Join to Source—Sending a join to the source. • Local RP—Sending neither join messages nor prune messages toward the RP, because this routing device is the rendezvous point. • Local Source—Sending neither join messages nor prune messages toward the source, because the source is locally attached to this routing device. • Prune to RP—Sending a prune to the rendezvous point. • Prune to Source—Sending a prune to the source. <p>NOTE: RP group range entries have None in the Upstream state field because RP group ranges do not trigger actual PIM join messages between routing devices.</p>	extensive
Downstream neighbors	<p>Information about downstream interfaces:</p> <ul style="list-style-type: none"> • Interface—Interface name for the downstream neighbor. A pseudo PIM-SM interface appears for all IGMP-only interfaces. A pseudo multipoint LDP (M-LDP) interface appears on ingress root nodes in M-LDP point-to-multipoint LSPs with inband signaling. • Interface address—Address of the downstream neighbor. • State—Information about the downstream neighbor: join or prune. • Flags—PIM join flags: R (RPtree), S (Sparse), W (Wildcard), or zero. • Uptime—Time since the downstream interface joined the group. • Time since last Join—Time since the last join message was received from the downstream interface. • Time since last Prune—Time since the last prune message was received from the downstream interface. 	extensive
Number of downstream interfaces	Total number of outgoing interfaces for each (S,G) entry.	extensive
Assert Timeout	Length of time between assert cycles on the downstream interface. Not displayed if the assert timer is null.	extensive

Table 67: show pim join Output Fields (*continued*)

Field Name	Field Description	Level of Output
Keepalive timeout	Time remaining until the downstream join state is updated (in seconds). If the downstream join state is not updated before this keepalive timer reaches zero, the entry is deleted. If there is a directly connected host, Keepalive timeout is Infinity .	extensive
Uptime	Time since the creation of (S,G) or (*,G) state. The uptime is not refreshed every time a PIM join message is received for an existing (S,G) or (*,G) state.	extensive
Bidirectional accepting interfaces	<p>Interfaces on the routing device that forward bidirectional PIM traffic.</p> <p>The reasons for forwarding bidirectional PIM traffic are that the interface is the winner of the designated forwarder election (DF Winner), or the interface is the reverse path forwarding (RPF) interface toward the RP (RPF).</p>	extensive

Sample Output

show pim join summary

```

user@host> show pim join summary
Instance: PIM.master Family: INET

Route type          Route count
(s,g)               2
(*,g)               1

Instance: PIM.master Family: INET6

```

show pim join (PIM Sparse Mode)

```

user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
  Source: *
  RP: 10.255.14.144
  Flags: sparse,rptree,wildcard
  Upstream interface: Local

Group: 239.1.1.1
  Source: 10.255.14.144
  Flags: sparse,spt
  Upstream interface: Local

Group: 239.1.1.1
  Source: 10.255.70.15
  Flags: sparse,spt
  Upstream interface: so-1/0/0.0

```

```
Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join (Bidirectional PIM)

```
user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.13.2
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0

Group: 224.1.3.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.1.3
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0 (RP Link)

Group: 225.1.1.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.13.2
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0

Group: 225.1.3.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.1.3
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0 (RP Link)

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join inet6

```
user@host> show pim join inet6
Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff04::e000:101
  Source: *
  RP: ::46.0.0.13
  Flags: sparse,rptree,wildcard
  Upstream interface: Local

Group: ff04::e000:101
  Source: ::1.1.1.1
  Flags: sparse
  Upstream interface: unknown (no neighbor)

Group: ff04::e800:101
  Source: ::1.1.1.1
  Flags: sparse
  Upstream interface: unknown (no neighbor)
```



```

Group: ff04::e800:101
Source: ::1.1.1.2
Flags: sparse
Upstream interface: unknown (no neighbor)

```

show pim join inet6 star-g

```

user@host> show pim join inet6 star-g
Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff04::e000:101
Source: *
RP: ::46.0.0.13
Flags: sparse,rptree,wildcard
Upstream interface: Local

```

show pim join instance <instance-name>

```

user@host> show pim join instance VPN-A
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 235.1.1.2
Source: *
RP: 10.10.47.100
Flags: sparse,rptree,wildcard
Upstream interface: Local

Group: 235.1.1.2
Source: 192.168.195.74
Flags: sparse,spt
Upstream interface: at-0/3/1.0

Group: 235.1.1.2
Source: 192.168.195.169
Flags: sparse
Upstream interface: so-1/0/1.0

Instance: PIM.VPN-A Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

show pim join detail

```

user@host> show pim join detail
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
Source: *
RP: 10.255.14.144
Flags: sparse,rptree,wildcard
Upstream interface: Local

Group: 239.1.1.1
Source: 10.255.14.144
Flags: sparse,spt
Upstream interface: Local

Group: 239.1.1.1
Source: 10.255.70.15

```

```
Flags: sparse,spt
Upstream interface: so-1/0/0.0
```

```
Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join extensive (PIM Sparse Mode)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
  Source: *
  RP: 10.255.14.144
  Flags: sparse,rptree,wildcard
  Upstream interface: Local
  Upstream neighbor: Local
  Upstream state: Local RP
  Uptime: 00:03:49
  Downstream neighbors:
    Interface: so-1/0/0.0
      10.111.10.2 State: Join Flags: SRW Timeout: 174
      Uptime: 00:03:49 Time since last Join: 00:01:49
    Interface: mt-1/1/0.32768
      10.10.47.100 State: Join Flags: SRW Timeout: Infinity
      Uptime: 00:03:49 Time since last Join: 00:01:49
  Number of downstream interfaces: 2

Group: 239.1.1.1
  Source: 10.255.14.144
  Flags: sparse,spt
  Upstream interface: Local
  Upstream neighbor: Local
  Upstream state: Local Source, Local RP
  Keepalive timeout: 344
  Uptime: 00:03:49
  Downstream neighbors:
    Interface: so-1/0/0.0
      10.111.10.2 State: Join Flags: S Timeout: 174
      Uptime: 00:03:49 Time since last Prune: 00:01:49
    Interface: mt-1/1/0.32768
      10.10.47.100 State: Join Flags: S Timeout: Infinity
      Uptime: 00:03:49 Time since last Prune: 00:01:49
  Number of downstream interfaces: 2

Group: 239.1.1.1
  Source: 10.255.70.15
  Flags: sparse,spt
  Upstream interface: so-1/0/0.0
  Upstream neighbor: 10.111.10.2
  Upstream state: Local RP, Join to Source
  Keepalive timeout: 344
  Uptime: 00:03:49
  Downstream neighbors:
    Interface: Pseudo-GMP
      fe-0/0/0.0 fe-0/0/1.0 fe-0/0/3.0
    Interface: so-1/0/0.0 (pruned)
      10.111.10.2 State: Prune Flags: SR Timeout: 174
      Uptime: 00:03:49 Time since last Prune: 00:01:49
    Interface: mt-1/1/0.32768
```

```

10.10.47.100 State: Join Flags: S   Timeout: Infinity
Uptime: 00:03:49 Time since last Prune: 00:01:49
Number of downstream interfaces: 3

```

```

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

show pim join extensive (Bidirectional PIM)

```

user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

```

Group: 224.1.1.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.13.2
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0
  Upstream neighbor: 10.10.1.2
  Upstream state: None
  Uptime: 00:03:49
  Bidirectional accepting interfaces:
    Interface: ge-0/0/1.0      (RPF)
    Interface: lo0.0          (DF Winner)
  Number of downstream interfaces: 0

Group: 225.1.1.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.13.2
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0
  Upstream neighbor: 10.10.1.2
  Upstream state: None
  Uptime: 00:03:49
  Bidirectional accepting interfaces:
    Interface: ge-0/0/1.0      (RPF)
    Interface: lo0.0          (DF Winner)
  Downstream neighbors:
    Interface: lt-1/0/10.24
      10.0.24.4 State: Join   RW   Timeout: 185
    Interface: lt-1/0/10.23
      10.0.23.3 State: Join   RW   Timeout: 184
  Number of downstream interfaces: 2

Group: 225.1.3.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.1.3
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0 (RP Link)
  Upstream neighbor: Direct
  Upstream state: Local RP
  Uptime: 00:03:49
  Bidirectional accepting interfaces:
    Interface: ge-0/0/1.0      (RPF)
    Interface: lo0.0          (DF Winner)
    Interface: xe-4/1/0.0      (DF Winner)
  Number of downstream interfaces: 0

```

```
Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join extensive (Bidirectional PIM with a Directly Connected Phantom RP)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.1
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.1.3
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0 (RP Link)
  Upstream neighbor: Direct
  Upstream state: Local RP
  Uptime: 00:03:49
  Bidirectional accepting interfaces:
    Interface: ge-0/0/1.0      (RPF)
    Interface: lo0.0          (DF Winner)
    Interface: xe-4/1/0.0     (DF Winner)
  Number of downstream interfaces: 0
```

show pim join instance <instance-name> extensive

```
user@host> show pim join instance VPN-A extensive
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 235.1.1.2
  Source: *
  RP: 10.10.47.100
  Flags: sparse,rptree,wildcard
  Upstream interface: Local
  Upstream neighbor: Local
  Upstream state: Local RP
  Uptime: 00:03:49
  Downstream neighbors:
    Interface: mt-1/1/0.32768
      10.10.47.101 State: Join Flags: SRW Timeout: 156
      Uptime: 00:03:49 Time since last Join: 00:01:49
  Number of downstream interfaces: 1

Group: 235.1.1.2
  Source: 192.168.195.74
  Flags: sparse,spt
  Upstream interface: at-0/3/1.0
  Upstream neighbor: 10.111.30.2
  Upstream state: Local RP, Join to Source
  Keepalive timeout: 156
  Uptime: 00:14:52

Group: 235.1.1.2
  Source: 192.168.195.169
  Flags: sparse
  Upstream interface: so-1/0/1.0
  Upstream neighbor: 10.111.20.2
  Upstream state: Local RP, Join to Source
  Keepalive timeout: 156
  Uptime: 00:14:52
```

show pim join extensive (Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```

user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 232.1.1.1
  Source: 192.168.219.11
  Flags: sparse,spt
  Upstream interface: fe-1/3/1.0
  Upstream neighbor: Direct
  Upstream state: Local Source
  Keepalive timeout:
  Uptime: 11:27:55
  Downstream neighbors:
    Interface: Pseudo-MLDP
    Interface: lt-1/2/0.25
      1.2.5.2 State: Join Flags: S   Timeout: Infinity
      Uptime: 11:27:55 Time since last Join: 11:27:55

Group: 232.1.1.2
  Source: 192.168.219.11
  Flags: sparse,spt
  Upstream interface: fe-1/3/1.0
  Upstream neighbor: Direct
  Upstream state: Local Source
  Keepalive timeout:
  Uptime: 11:27:41
  Downstream neighbors:
    Interface: Pseudo-MLDP

Group: 232.1.1.3
  Source: 192.168.219.11
  Flags: sparse,spt
  Upstream interface: fe-1/3/1.0
  Upstream neighbor: Direct
  Upstream state: Local Source
  Keepalive timeout:
  Uptime: 11:27:41
  Downstream neighbors:
    Interface: Pseudo-MLDP

Group: 232.2.2.2
  Source: 1.2.7.7
  Flags: sparse,spt
  Upstream interface: lt-1/2/0.27
  Upstream neighbor: Direct
  Upstream state: Local Source
  Keepalive timeout:
  Uptime: 11:27:25
  Downstream neighbors:
    Interface: Pseudo-MLDP

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff3e::1:2
  Source: abcd::1:2:7:7
  Flags: sparse,spt
  Upstream interface: lt-1/2/0.27
  Upstream neighbor: Direct

```

```
Upstream state: Local Source
Keepalive timeout:
Uptime: 11:27:26
Downstream neighbors:
  Interface: Pseudo-MLDP
```

show pim join extensive (Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 227.1.1.1
  Source: *
  RP: 1.1.1.1
  Flags: sparse,rptree,wildcard
  Upstream interface: Local
  Upstream neighbor: Local
  Upstream state: Local RP
  Uptime: 11:31:33
  Downstream neighbors:
    Interface: fe-1/3/0.0
      192.168.209.9 State: Join Flags: SRW Timeout: Infinity
      Uptime: 11:31:33 Time since last Join: 11:31:32

Group: 232.1.1.1
  Source: 192.168.219.11
  Flags: sparse,spt
  Upstream protocol: MLDP
  Upstream interface: Pseudo MLDP
  Upstream neighbor: MLDP LSP root <1.1.1.2>
  Upstream state: Join to Source
  Keepalive timeout:
  Uptime: 11:31:32
  Downstream neighbors:
    Interface: so-0/1/3.0
      192.168.92.9 State: Join Flags: S Timeout: Infinity
      Uptime: 11:31:30 Time since last Join: 11:31:30
    Downstream neighbors:
      Interface: fe-1/3/0.0
        192.168.209.9 State: Join Flags: S Timeout: Infinity
        Uptime: 11:31:32 Time since last Join: 11:31:32

Group: 232.1.1.2
  Source: 192.168.219.11
  Flags: sparse,spt
  Upstream protocol: MLDP
  Upstream interface: Pseudo MLDP
  Upstream neighbor: MLDP LSP root <1.1.1.2>
  Upstream state: Join to Source
  Keepalive timeout:
  Uptime: 11:31:32
  Downstream neighbors:
    Interface: so-0/1/3.0
      192.168.92.9 State: Join Flags: S Timeout: Infinity
      Uptime: 11:31:30 Time since last Join: 11:31:30
    Downstream neighbors:
      Interface: lt-1/2/0.14
        1.1.4.4 State: Join Flags: S Timeout: 177
        Uptime: 11:30:33 Time since last Join: 00:00:33
    Downstream neighbors:
```

```

Interface: fe-1/3/0.0
  192.168.209.9 State: Join Flags: S   Timeout: Infinity
  Uptime: 11:31:32 Time since last Join: 11:31:32

Group: 232.1.1.3
  Source: 192.168.219.11
  Flags: sparse,spt
  Upstream protocol: MLDP
  Upstream interface: Pseudo MLDP
  Upstream neighbor: MLDP LSP root <1.1.1.2>
  Upstream state: Join to Source
  Keepalive timeout:
  Uptime: 11:31:32
  Downstream neighbors:
    Interface: fe-1/3/0.0
      192.168.209.9 State: Join Flags: S   Timeout: Infinity
      Uptime: 11:31:32 Time since last Join: 11:31:32

Group: 232.2.2.2
  Source: 1.2.7.7
  Flags: sparse,spt
  Upstream protocol: MLDP
  Upstream interface: Pseudo MLDP
  Upstream neighbor: MLDP LSP root <1.1.1.2>
  Upstream state: Join to Source
  Keepalive timeout:
  Uptime: 11:31:30
  Downstream neighbors:
    Interface: so-0/1/3.0
      192.168.92.9 State: Join Flags: S   Timeout: Infinity
      Uptime: 11:31:30 Time since last Join: 11:31:30

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff3e::1:2
  Source: abcd::1:2:7:7
  Flags: sparse,spt
  Upstream protocol: MLDP
  Upstream interface: Pseudo MLDP
  Upstream neighbor: MLDP LSP root <1.1.1.2>
  Upstream state: Join to Source
  Keepalive timeout:
  Uptime: 11:31:32
  Downstream neighbors:
    Interface: fe-1/3/0.0
      fe80::21f:12ff:fea5:c4db State: Join Flags: S   Timeout: Infinity
      Uptime: 11:31:32 Time since last Join: 11:31:32

```

Sample Output

show pim join summary

```

user@host> show pim join summary
Instance: PIM.master Family: INET

Route type      Route count
(s,g)           2
(*,g)           1

Instance: PIM.master Family: INET6

```

show pim join (PIM Sparse Mode)

```
user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
Source: *
RP: 10.255.14.144
Flags: sparse,rptree,wildcard
Upstream interface: Local

Group: 239.1.1.1
Source: 10.255.14.144
Flags: sparse,spt
Upstream interface: Local

Group: 239.1.1.1
Source: 10.255.70.15
Flags: sparse,spt
Upstream interface: so-1/0/0.0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join (Bidirectional PIM)

```
user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.0
Bidirectional group prefix length: 24
Source: *
RP: 10.10.13.2
Flags: bidirectional,rptree,wildcard
Upstream interface: ge-0/0/1.0

Group: 224.1.3.0
Bidirectional group prefix length: 24
Source: *
RP: 10.10.1.3
Flags: bidirectional,rptree,wildcard
Upstream interface: ge-0/0/1.0 (RP Link)

Group: 225.1.1.0
Bidirectional group prefix length: 24
Source: *
RP: 10.10.13.2
Flags: bidirectional,rptree,wildcard
Upstream interface: ge-0/0/1.0

Group: 225.1.3.0
Bidirectional group prefix length: 24
Source: *
RP: 10.10.1.3
Flags: bidirectional,rptree,wildcard
Upstream interface: ge-0/0/1.0 (RP Link)

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```


show pim join inet6

```

user@host> show pim join inet6
Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff04::e000:101
  Source: *
  RP: ::46.0.0.13
  Flags: sparse,rptree,wildcard
  Upstream interface: Local

Group: ff04::e000:101
  Source: ::1.1.1.1
  Flags: sparse
  Upstream interface: unknown (no neighbor)

Group: ff04::e800:101
  Source: ::1.1.1.1
  Flags: sparse
  Upstream interface: unknown (no neighbor)

Group: ff04::e800:101
  Source: ::1.1.1.2
  Flags: sparse
  Upstream interface: unknown (no neighbor)

```

show pim join inet6 star-g

```

user@host> show pim join inet6 star-g
Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff04::e000:101
  Source: *
  RP: ::46.0.0.13
  Flags: sparse,rptree,wildcard
  Upstream interface: Local

```

show pim join instance <instance-name>

```

user@host> show pim join instance VPN-A
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 235.1.1.2
  Source: *
  RP: 10.10.47.100
  Flags: sparse,rptree,wildcard
  Upstream interface: Local

Group: 235.1.1.2
  Source: 192.168.195.74
  Flags: sparse,spt
  Upstream interface: at-0/3/1.0

Group: 235.1.1.2
  Source: 192.168.195.169
  Flags: sparse
  Upstream interface: so-1/0/1.0

```

```
Instance: PIM.VPN-A Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join detail

```
user@host> show pim join detail
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
  Source: *
  RP: 10.255.14.144
  Flags: sparse,rptree,wildcard
  Upstream interface: Local

Group: 239.1.1.1
  Source: 10.255.14.144
  Flags: sparse,spt
  Upstream interface: Local

Group: 239.1.1.1
  Source: 10.255.70.15
  Flags: sparse,spt
  Upstream interface: so-1/0/0.0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join extensive (PIM Sparse Mode)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
  Source: *
  RP: 10.255.14.144
  Flags: sparse,rptree,wildcard
  Upstream interface: Local
  Upstream neighbor: Local
  Upstream state: Local RP
  Uptime: 00:03:49
  Downstream neighbors:
    Interface: so-1/0/0.0
      10.111.10.2 State: Join Flags: SRW Timeout: 174
      Uptime: 00:03:49 Time since last Join: 00:01:49
    Interface: mt-1/1/0.32768
      10.10.47.100 State: Join Flags: SRW Timeout: Infinity
      Uptime: 00:03:49 Time since last Join: 00:01:49
  Number of downstream interfaces: 2

Group: 239.1.1.1
  Source: 10.255.14.144
  Flags: sparse,spt
  Upstream interface: Local
  Upstream neighbor: Local
  Upstream state: Local Source, Local RP
  Keepalive timeout: 344
  Uptime: 00:03:49
  Downstream neighbors:
    Interface: so-1/0/0.0
```

```

    10.111.10.2 State: Join Flags: S Timeout: 174
    Uptime: 00:03:49 Time since last Prune: 00:01:49
    Interface: mt-1/1/0.32768
    10.10.47.100 State: Join Flags: S   Timeout: Infinity
    Uptime: 00:03:49 Time since last Prune: 00:01:49
    Number of downstream interfaces: 2

Group: 239.1.1.1
Source: 10.255.70.15
Flags: sparse,spt
Upstream interface: so-1/0/0.0
Upstream neighbor: 10.111.10.2
Upstream state: Local RP, Join to Source
Keepalive timeout: 344
Uptime: 00:03:49
Downstream neighbors:
  Interface: Pseudo-GMP
    fe-0/0/0.0 fe-0/0/1.0 fe-0/0/3.0
  Interface: so-1/0/0.0 (pruned)
    10.111.10.2 State: Prune Flags: SR Timeout: 174
    Uptime: 00:03:49 Time since last Prune: 00:01:49
  Interface: mt-1/1/0.32768
    10.10.47.100 State: Join Flags: S   Timeout: Infinity
    Uptime: 00:03:49 Time since last Prune: 00:01:49
  Number of downstream interfaces: 3

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

```

show pim join extensive (Bidirectional PIM)

```

user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.1.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.13.2
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0
  Upstream neighbor: 10.10.1.2
  Upstream state: None
  Uptime: 00:03:49
  Bidirectional accepting interfaces:
    Interface: ge-0/0/1.0      (RPF)
    Interface: lo0.0           (DF Winner)
  Number of downstream interfaces: 0

Group: 225.1.1.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.13.2
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0
  Upstream neighbor: 10.10.1.2
  Upstream state: None
  Uptime: 00:03:49
  Bidirectional accepting interfaces:
    Interface: ge-0/0/1.0      (RPF)
    Interface: lo0.0           (DF Winner)

```

```
Downstream neighbors:
  Interface: lt-1/0/10.24
    10.0.24.4 State: Join   RW   Timeout: 185
  Interface: lt-1/0/10.23
    10.0.23.3 State: Join   RW   Timeout: 184
  Number of downstream interfaces: 2

Group: 225.1.3.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.1.3
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0 (RP Link)
  Upstream neighbor: Direct
  Upstream state: Local RP
  Uptime: 00:03:49
  Bidirectional accepting interfaces:
    Interface: ge-0/0/1.0      (RPF)
    Interface: lo0.0           (DF Winner)
    Interface: xe-4/1/0.0      (DF Winner)
  Number of downstream interfaces: 0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join extensive (Bidirectional PIM with a Directly Connected Phantom RP)

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 224.1.3.0
  Bidirectional group prefix length: 24
  Source: *
  RP: 10.10.1.3
  Flags: bidirectional,rptree,wildcard
  Upstream interface: ge-0/0/1.0 (RP Link)
  Upstream neighbor: Direct
  Upstream state: Local RP
  Uptime: 00:03:49
  Bidirectional accepting interfaces:
    Interface: ge-0/0/1.0      (RPF)
    Interface: lo0.0           (DF Winner)
    Interface: xe-4/1/0.0      (DF Winner)
  Number of downstream interfaces: 0
```

show pim join instance <instance-name> extensive

```
user@host> show pim join instance VPN-A extensive
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 235.1.1.2
  Source: *
  RP: 10.10.47.100
  Flags: sparse,rptree,wildcard
  Upstream interface: Local
  Upstream neighbor: Local
  Upstream state: Local RP
  Uptime: 00:03:49
  Downstream neighbors:
```

```

Interface: mt-1/1/0.32768
10.10.47.101 State: Join Flags: SRW Timeout: 156
Uptime: 00:03:49 Time since last Join: 00:01:49
Number of downstream interfaces: 1

Group: 235.1.1.2
Source: 192.168.195.74
Flags: sparse,spt
Upstream interface: at-0/3/1.0
Upstream neighbor: 10.111.30.2
Upstream state: Local RP, Join to Source
Keepalive timeout: 156
Uptime: 00:14:52

Group: 235.1.1.2
Source: 192.168.195.169
Flags: sparse
Upstream interface: so-1/0/1.0
Upstream neighbor: 10.111.20.2
Upstream state: Local RP, Join to Source
Keepalive timeout: 156
Uptime: 00:14:52

```

show pim join extensive (Ingress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```

user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 232.1.1.1
Source: 192.168.219.11
Flags: sparse,spt
Upstream interface: fe-1/3/1.0
Upstream neighbor: Direct
Upstream state: Local Source
Keepalive timeout:
Uptime: 11:27:55
Downstream neighbors:
Interface: Pseudo-MLDP
Interface: lt-1/2/0.25
1.2.5.2 State: Join Flags: S Timeout: Infinity
Uptime: 11:27:55 Time since last Join: 11:27:55

Group: 232.1.1.2
Source: 192.168.219.11
Flags: sparse,spt
Upstream interface: fe-1/3/1.0
Upstream neighbor: Direct
Upstream state: Local Source
Keepalive timeout:
Uptime: 11:27:41
Downstream neighbors:
Interface: Pseudo-MLDP

Group: 232.1.1.3
Source: 192.168.219.11
Flags: sparse,spt
Upstream interface: fe-1/3/1.0
Upstream neighbor: Direct
Upstream state: Local Source
Keepalive timeout:

```

```
Uptime: 11:27:41
Downstream neighbors:
  Interface: Pseudo-MLDP

Group: 232.2.2.2
Source: 1.2.7.7
Flags: sparse,spt
Upstream interface: lt-1/2/0.27
Upstream neighbor: Direct
Upstream state: Local Source
Keepalive timeout:
Uptime: 11:27:25
Downstream neighbors:
  Interface: Pseudo-MLDP

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: ff3e::1:2
Source: abcd::1:2:7:7
Flags: sparse,spt
Upstream interface: lt-1/2/0.27
Upstream neighbor: Direct
Upstream state: Local Source
Keepalive timeout:
Uptime: 11:27:26
Downstream neighbors:
  Interface: Pseudo-MLDP
```

show pim join extensive (Multipoint LDP with Multicast-Only Fast Reroute)

```
user@host> show pim join 225.1.1.1 extensive sg
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 225.1.1.1
Source: 10.0.0.1
Flags: sparse,spt
Active upstream interface: fe-1/2/13.0
Active upstream neighbor: 10.0.0.9
MoFRR Backup upstream interface: fe-1/2/14.0
MoFRR Backup upstream neighbor: 10.0.0.21
Upstream state: Join to Source, No Prune to RP
Keepalive timeout: 354
Uptime: 00:00:06
Downstream neighbors:
  Interface: fe-1/2/15.0
    10.0.0.13 State: Join Flags: S   Timeout: Infinity
    Uptime: 00:00:06 Time since last Join: 00:00:06
Number of downstream interfaces: 1
```

show pim mdt

Syntax	show pim mdt instance <i>instance-name</i> <brief detail extensive> <incoming outgoing> <logical-system (all logical-system-name)> <range>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display information about Protocol Independent Multicast (PIM) default multicast distribution tree (MDT) and the data MDTs in a Layer 3 VPN environment for a routing instance.
Options	<p>instance <i>instance-name</i>—Display information about data-MDTs for a specific PIM-enabled routing instance.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>incoming outgoing—(Optional) Display incoming or outgoing multicast data tunnels, respectively.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>range—(Optional) Display information about an IP address with optional prefix length representing a particular multicast group.</p>
Required Privilege Level	view
List of Sample Output	show pim mdt instance on page 534 show pim mdt instance detail on page 535 show pim mdt instance extensive on page 535 show pim mdt instance incoming on page 535 show pim mdt instance outgoing on page 535 show pim mdt instance (SSM Mode) on page 536
Output Fields	Table 68 on page 533 describes the output fields for the show pim mdt command. Output fields are listed in the approximate order in which they appear.

Table 68: show pim mdt Output Fields

Field Name	Field Description	Level of Output
Instance	Name of the routing instance.	All levels
Tunnel direction	Direction the tunnel faces, from the router's perspective: Outgoing or Incoming .	All levels
Tunnel mode	Mode the tunnel is operating in: PIM-SSM or PIM-ASM .	All levels

Table 68: show pim mdt Output Fields (*continued*)

Field Name	Field Description	Level of Output
Default group address	Default multicast group address using this tunnel.	All levels
Default source address	Default multicast source address using this tunnel.	All levels
Default tunnel interface	Default multicast tunnel interface.	All levels
Default tunnel source	Address used as the source address for outgoing PIM control messages.	All levels
C-Group	Customer-facing multicast group address using this tunnel. If you enable dynamic reuse of data MDT group addresses, more than one group address can use the same data MDT.	detail
C-Source	IP address of the multicast source in the customer's address space. If you enable dynamic reuse of data MDT group addresses, more than one source address can use the same data MDT.	detail
P-Group	Service provider-facing multicast group address using this tunnel.	detail
Data tunnel interface	Multicast data tunnel interface that set up the data-MDT tunnel.	detail
Last known forwarding rate	Last known rate, in kilobits per second, at which the tunnel was forwarding traffic.	detail
Configured threshold rate	Rate, in kilobits per second, above which a data-MDT tunnel is created and below which it is deleted.	detail
Tunnel uptime	Time that this data-MDT tunnel has existed. The format is <i>hours:minutes:seconds</i> .	detail

Sample Output

show pim mdt instance

```
user@host> show pim mdt instance VPN-A
```

```
Instance: PIM.VPN-A
```

```
Tunnel direction: Outgoing
```

```
Default group address: 239.1.1.1
```

```
Default tunnel interface: mt-1/1/0.32768
```

```
Default tunnel source: 192.168.7.1
```

```

C-group address  C-source address  P-group address  Data tunnel interface
235.1.1.2        192.168.195.74    228.0.0.0        mt-1/1/0.32769

```

```
Instance: PIM.VPN-A
```

```
Tunnel direction: Incoming
```

```
Default group address: 239.1.1.1
```

```
Default tunnel interface: mt-1/1/0.1081344
```


show pim mdt instance detail

```

user@host> show pim mdt instance VPN-A detail
Instance: PIM.VPN-A
Tunnel direction: Outgoing
Default group address: 239.1.1.1
Default tunnel interface: mt-1/1/0.32768
Default tunnel source: 192.168.7.1

C-Group: 235.1.1.2
  C-Source: 192.168.195.74
  P-Group : 228.0.0.0
  Data tunnel interface      : mt-1/1/0.32769
  Last known forwarding rate : 48 kbps (6 kbps)
  Configured threshold rate  : 10 kbps
  Tunnel uptime              : 00:00:34

Instance: PIM.VPN-A
Tunnel direction: Incoming
Default group address: 239.1.1.1
Default tunnel interface: mt-1/1/0.1081344

```

show pim mdt instance extensive

```

user@host> show pim mdt instance VPN-A extensive
Instance: PIM.VPN-A
Tunnel direction: Outgoing
Default group address: 239.1.1.1
Default tunnel interface: mt-1/1/0.32768
Default tunnel source: 192.168.7.1

C-Group: 235.1.1.2
  C-Source: 192.168.195.74
  P-Group : 228.0.0.0
  Data tunnel interface      : mt-1/1/0.32769
  Last known forwarding rate : 48 kbps (6 kbps)
  Configured threshold rate  : 10 kbps
  Tunnel uptime              : 00:00:41

Instance: PIM.VPN-A
Tunnel direction: Incoming
Default group address: 239.1.1.1
Default tunnel interface: mt-1/1/0.1081344

```

show pim mdt instance incoming

```

user@host> show pim mdt instance VPN-A incoming
Instance: PIM.VPN-A
Tunnel direction: Incoming
Default group address: 239.1.1.1
Default tunnel interface: mt-1/1/0.1081344

```

show pim mdt instance outgoing

```

user@host> show pim mdt instance VPN-A outgoing
Instance: PIM.VPN-A
Tunnel direction: Outgoing
Default group address: 239.1.1.1
Default tunnel interface: mt-1/1/0.32768
Default tunnel source: 192.168.7.1

```

C-group address	C-source address	P-group address	Data tunnel interface
235.1.1.2	192.168.195.74	228.0.0.0	mt-1/1/0.32769

show pim mdt instance (SSM Mode)

```
user@host> show pim mdt instance vpn-a
Instance: PIM.vpn-a
Tunnel direction: Outgoing
Tunnel mode: PIM-SSM
Default group address: 232.1.1.1
Default source address: 10.255.14.216
Default tunnel interface: mt-1/3/0.32769
Default tunnel source: 192.168.7.1

Instance: PIM.vpn-a
Tunnel direction: Incoming
Tunnel mode: PIM-SSM
Default group address: 232.1.1.1
Default source address: 10.255.14.217
Default tunnel interface: mt-1/3/0.1081345

Instance: PIM.vpn-a
Tunnel direction: Incoming
Tunnel mode: PIM-SSM
Default group address: 232.1.1.1
Default source address: 10.255.14.218
Default tunnel interface: mt-1/3/0.1081345
```

show pim mdt data-mdt-joins

Syntax `show pim mdt data-mdt-joins`
`<logical-system (all | logical-system-name)> instance instance-name`

Release Information Command introduced in Junos OS Release 11.2.

Description In a draft-rosen Layer 3 multicast virtual private network (MVPN) configured with service provider tunnels, display the advertisements of new multicast distribution tree (MDT) group addresses cached by the provider edge (PE) routers in the specified VPN routing and forwarding (VRF) instance that is configured to use the Protocol Independent Multicast (PIM) protocol.

Options `instance instance-name`—Display data MDT join packets cached by PE routers in a specific PIM instance.

`logical-system (all | logical-system-name)`—(Optional) Perform this operation on all logical systems or on a particular logical system.



NOTE: Draft-rosen multicast VPNs are not supported in a logical system environment even though the configuration statements can be configured under the logical-systems hierarchy.

Required Privilege Level view

Related Documentation

- *Understanding Data MDTs*
- *Example: Configuring Data MDTs and Provider Tunnels Operating in Source-Specific Multicast Mode*
- *Example: Configuring Data MDTs and Provider Tunnels Operating in Any-Source Multicast Mode*

List of Sample Output [show pim mdt data-mdt-joins on page 538](#)

Output Fields [Table 69 on page 537](#) describes the output fields for the `show pim mdt data-mdt-joins` command. Output fields are listed in the approximate order in which they appear.

Table 69: show pim mdt data-mdt-joins Output Fields

Field Name	Field Description
C-Group	IPv4 group address in the address space of the customer's VPN-specific PIM-enabled routing instance of the multicast traffic destination. This 32-bit value is carried in the C-group field of the MDT join TLV packet.
C-Source	IPv4 address in the address space of the customer's VPN-specific PIM-enabled routing instance of the multicast traffic source. This 32-bit value is carried in the C-source field of the MDT join TLV packet.

Table 69: show pim mdt data-mdt-joins Output Fields (*continued*)

Field Name	Field Description
P-Group	IPv4 group address in the service provider's address space of the new data MDT that the PE router will use to encapsulate the VPN multicast traffic flow (C-Source, C-Group). This 32-bit value is carried in the P-group field of the MDT join TLV packet.
P-Source	IPv4 address of the PE router.
Timeout	Timeout, in seconds, remaining for this cache entry. When the cache entry is created, this field is set to 180 seconds. After an entry times out, the PE router deletes the entry from its cache and prunes itself off the data MDT.

Sample Output

show pim mdt data-mdt-joins

```

user@host show pim mdt data-mdt-joins instance VPN-A
C-Source      C-Group      P-Source      P-Group      Timeout
20.2.15.9     225.1.1.2    20.0.0.5      239.10.10.0  172
20.2.15.9     225.1.1.3    20.0.0.5      239.10.10.1  172

```

show pim mdt data-mdt-limit

Syntax `show pim mdt data-mdt-limit instance instance-name`
`<logical-system (all | logical-system-name)>`

Release Information Command introduced in Junos OS Release 12.2.

Description Display the maximum number configured and the currently active data multicast distribution trees (MDTs) for a specific VPN routing and forwarding (VRF) instance.

Options **instance *instance-name***—Display data MDT information for the specified VRF instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.



NOTE: Draft-rosen multicast VPNs are not supported in a logical system environment even though the configuration statements can be configured under the logical-systems hierarchy.

Required Privilege Level view

Related Documentation

- *Understanding Data MDTs*
- *Example: Configuring Data MDTs and Provider Tunnels Operating in Source-Specific Multicast Mode*
- *Example: Configuring Data MDTs and Provider Tunnels Operating in Any-Source Multicast Mode*

List of Sample Output [show pim mdt data-mdt-limit on page 539](#)

Output Fields [Table 70 on page 539](#) describes the output fields for the **show pim mdt data-mdt-limit** command. Output fields are listed in the approximate order in which they appear.

Table 70: show pim mdt data-mdt-limit Output Fields

Field Name	Field Description
Maximum Data Tunnels	Maximum number of data MDTs created in this VRF instance. If the number is 0, no data MDTs are created for this VRF instance.
Active Data Tunnels	Active number of data MDTs in this VRF instance.

Sample Output

show pim mdt data-mdt-limit

```
user@host show pim mdt data-mdt-limit instance VPN-A
```

Maximum Data Tunnels	10
Active Data Tunnels	2

show pim neighbors

Syntax	<pre>show pim neighbors <brief detail> <inet inet6> <instance (instance-name all)> <logical-system (all logical-system-name)></pre>
Syntax (EX Series Switch and the QFX Series)	<pre>show pim neighbors <brief detail> <inet inet6> <instance (instance-name all)></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p> <p>Support for the instance all option added in Junos OS Release 12.1.</p>
Description	Display information about Protocol Independent Multicast (PIM) neighbors.
Options	<p>none—(Same as brief) Display standard information about PIM neighbors for all supported family addresses for the main instance.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>inet inet6—(Optional) Display information about PIM neighbors for IPv4 or IPv6 family addresses, respectively.</p> <p>instance (instance-name all)—(Optional) Display information about neighbors for the specified PIM-enabled routing instance or for all routing instances.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	<p>show pim neighbors on page 543</p> <p>show pim neighbors brief on page 543</p> <p>show pim neighbors instance on page 543</p> <p>show pim neighbors detail on page 543</p> <p>show pim neighbors detail (With BFD) on page 544</p>
Output Fields	<p>Table 71 on page 542 describes the output fields for the show pim neighbors command. Output fields are listed in the approximate order in which they appear.</p>

Table 71: show pim neighbors Output Fields

Field Name	Field Description	Level of Output
Instance	Name of the routing instance.	All levels
Interface	Interface through which the neighbor is reachable.	All levels
Neighbor addr	Address of the neighboring PIM routing device.	All levels
IP	IP version: 4 or 6.	All levels
V	PIM version running on the neighbor: 1 or 2.	All levels
Mode	PIM mode of the neighbor: Sparse , Dense , SparseDense , or Unknown . When the neighbor is running PIM version 2, this mode is always Unknown .	All levels
Option	Can be one or more of the following: <ul style="list-style-type: none"> • B—Bidirectional Capable. • H—Hello Option Holdtime. • G—Generation Identifier. • P—Hello Option DR Priority. • L—Hello Option LAN Prune Delay. 	brief none
Uptime	Time the neighbor has been operational since the PIM process was last initialized, in the format dd:hh:mm:ss ago for less than a week and nwnd:hh:mm:ss ago for more than a week.	All levels
Address	Address of the neighboring PIM routing device.	detail
BFD	Status and operational state of the Bidirectional Forwarding Detection (BFD) protocol on the interface: Enabled , Operational state is up , or Disabled .	detail
Hello Option Holdtime	Time for which the neighbor is available, in seconds. The range of values is 0 through 65,535.	detail
Hello Default Holdtime	Default holdtime and the time remaining if the holdtime option is not in the received hello message.	detail
Hello Option DR Priority	Designated router election priority. The range of values is 0 through 255.	detail
Hello Option Generation ID	9-digit or 10-digit number used to tag hello messages.	detail
Hello Option Bi-Directional PIM supported	Neighbor can process bidirectional PIM messages.	detail
Hello Option LAN Prune Delay	Time to wait before the neighbor receives prune messages, in the format delay nnn ms override nnnn ms .	detail

Table 71: show pim neighbors Output Fields (*continued*)

Field Name	Field Description	Level of Output
Join Suppression supported	Neighbor is capable of join suppression.	detail
Rx Join	Information about joins received from the neighbor. <ul style="list-style-type: none"> Group—Group addresses in the join message. Source—Address of the source in the join message. Timeout—Time for which the join is valid. 	detail

Sample Output

show pim neighbors

```

user@host> show pim neighbors
Instance: PIM.master
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority

Interface      IP V Mode      Option      Uptime Neighbor addr
so-1/0/0.0      4 2            HPLG        00:07:10 10.111.10.2

```

show pim neighbors brief

The output for the **show pim neighbors brief** command is identical to that for the **show pim neighbors** command. For sample output, see [show pim neighbors on page 543](#).

show pim neighbors instance

```

user@host> show pim neighbors instance VPN-A
Instance: PIM.VPN-A
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority

Interface      IP V Mode      Option      Uptime Neighbor addr
at-0/3/1.0      4 2            HPLG        00:07:54 10.111.30.2
mt-1/1/0.32768  4 2            HPLG        00:07:22 10.10.47.101
so-1/0/1.0      4 2            HPLG        00:07:50 10.111.20.2

```

show pim neighbors detail

```

user@host> show pim neighbors detail
Instance: PIM.master
Interface: ge-0/0/1.0

Address: 10.10.1.1, IPv4, PIM v2, Mode: SparseDense, sg Join Count: 0, ts
Join Count: 2
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 1
Hello Option Generation ID: 2053759302
Hello Option Bi-Directional PIM supported
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported

```

```
Address: 10.10.1.2, IPv4, PIM v2, sg Join Count: 0, tsg Join Count: 2
BFD: Disabled
Hello Option Holdtime: 105 seconds 93 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 1734018161
Hello Option Bi-Directional PIM supported
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported
```

Interface: lo0.0

```
Address: 10.255.179.246, IPv4, PIM v2, Mode: SparseDense, sg Join Count:
0, tsg Join Count: 0
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 1
Hello Option Generation ID: 1997462267
Hello Option Bi-Directional PIM supported
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Join Suppression supported
```

show pim neighbors detail (With BFD)

```
user@host> show pim neighbors detail
```

Instance: PIM.master

Interface: fe-1/0/0.0

```
Address: 192.168.11.1, IPv4, PIM v2, Mode: Sparse
Hello Option Holdtime: 65535 seconds
Hello Option DR Priority: 1
Hello Option Generation ID: 836607909
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

```
Address: 192.168.11.2, IPv4, PIM v2
BFD: Enabled, Operational state is up
Hello Default Holdtime: 105 seconds 104 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 1907549685
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

Interface: fe-1/0/1.0

```
Address: 192.168.12.1, IPv4, PIM v2
BFD: Disabled
Hello Default Holdtime: 105 seconds 80 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 1971554705
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

show pim rps

Syntax	<pre>show pim rps <brief detail extensive> <group-address> <inet inet6> <instance instance-name> <logical-system (all logical-system-name)></pre>
Syntax (EX Series Switch and the QFX Series)	<pre>show pim rps <brief detail extensive> <group-address> <inet inet6> <instance instance-name></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p> <p>Support for bidirectional PIM added in Junos OS Release 12.1.</p>
Description	Display information about Protocol Independent Multicast (PIM) rendezvous points (RPs).
Options	<p>none—Display standard information about PIM RPs for all groups and family addresses for all routing instances.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>group-address—(Optional) Display the RPs for a particular group. If you specify a group address, the output lists the routing device that is the RP for that group.</p> <p>inet inet6—(Optional) Display information for IPv4 or IPv6 family addresses, respectively.</p> <p>instance instance-name—(Optional) Display information about RPs for a specific PIM-enabled routing instance.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Bidirectional PIM</i>
List of Sample Output	<p>show pim rps on page 548</p> <p>show pim rps brief on page 548</p> <p>show pim rps <group-address> (Bidirectional PIM) on page 548</p> <p>show pim rps <group-address> (PIM Dense Mode) on page 548</p>

[show pim rps <group-address> \(SSM Range Without asm-override-ssm Configured\) on page 548](#)
[show pim rps <group-address> \(SSM Range With asm-override-ssm Configured and a Sparse-Mode RP\) on page 549](#)
[show pim rps <group-address> \(SSM Range With asm-override-ssm Configured and a Bidirectional RP\) on page 549](#)
[show pim rps instance on page 549](#)
[show pim rps extensive \(PIM Sparse Mode\) on page 549](#)
[show pim rps extensive \(Bidirectional PIM\) on page 550](#)
[show pim rps extensive \(PIM Anycast RP in Use\) on page 550](#)

Output Fields [Table 72 on page 546](#) describes the output fields for the **show pim rps** command. Output fields are listed in the approximate order in which they appear.

Table 72: show pim rps Output Fields

Field Name	Field Description	Level of Output
Instance	Name of the routing instance.	All levels
Family or Address family	Name of the address family: inet (IPv4) or inet6 (IPv6).	All levels
RP address	Address of the rendezvous point.	All levels
Type	Type of RP: <ul style="list-style-type: none"> auto-rp—Address of the RP known through the Auto-RP protocol. bootstrap—Address of the RP known through the bootstrap router protocol (BSR). embedded—Address of the RP known through an embedded RP (IPv6). static—Address of RP known through static configuration. 	brief none
Holdtime	How long to keep the RP active, with time remaining, in seconds.	All levels
Timeout	How long until the local routing device determines the RP to be unreachable, in seconds.	All levels
Groups	Number of groups currently using this RP.	All levels
Group prefixes	Addresses of groups that this RP can span.	brief none
Learned via	Address and method by which the RP was learned.	detail extensive
Mode	The PIM mode of the RP: bidirectional or sparse. If a sparse and bidirectional RPs are configured with the same RP address, they appear as separate entries in both formats.	All levels
Time Active	How long the RP has been active, in the format <i>hh:mm:ss</i> .	detail extensive

Table 72: show pim rps Output Fields (*continued*)

Field Name	Field Description	Level of Output
Device Index	Index value of the order in which Junos OS finds and initializes the interface. For bidirectional RPs, the Device Index output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces.	detail extensive
Subunit	Logical unit number of the interface. For bidirectional RPs, the Subunit output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces.	detail extensive
Interface	Either the encapsulation or the de-encapsulation logical interface, depending on whether this routing device is a designated router (DR) facing an RP router, or is the local RP, respectively. For bidirectional RPs, the Interface output field is omitted because bidirectional RPs do not require encapsulation and de-encapsulation interfaces.	detail extensive
Group Ranges	Addresses of groups that this RP spans.	detail extensive <i>group-address</i>
Active groups using RP	Number of groups currently using this RP.	detail extensive
total	Total number of active groups for this RP.	detail extensive
Register State for RP	Current register state for each group: <ul style="list-style-type: none"> • Group—Multicast group address. • Source—Multicast source address for which the PIM register is sent or received, depending on whether this router is a designated router facing an RP router, or is the local RP, respectively: • First Hop—PIM-designated routing device that sent the Register message (the source address in the IP header). • RP Address—RP to which the Register message was sent (the destination address in the IP header). • State: On the designated router: <ul style="list-style-type: none"> • Send—Sending Register messages. • Probe—Sent a null register. If a Register-Stop message does not arrive in 5 seconds, the designated router resumes sending Register messages. • Suppress—Received a Register-Stop message. The designated router is waiting for the timer to resume before changing to Probe state. • On the RP: <ul style="list-style-type: none"> • Receive—Receiving Register messages. 	extensive
Anycast-PIM rpset	If anycast RP is configured, the addresses of the RPs in the set.	extensive
Anycast-PIM local address used	If anycast RP is configured, the local address used by the RP.	extensive

Table 72: show pim rps Output Fields (*continued*)

Field Name	Field Description	Level of Output
Anycast-PIM Register State	<p>If anycast RP is configured, the current register state for each group:</p> <ul style="list-style-type: none"> • Group—Multicast group address. • Source—Multicast source address for which the PIM register is sent or received, depending on whether this routing device is a designated router facing an RP router, or is the local RP, respectively. • Origin—How the information was obtained: <ul style="list-style-type: none"> • DIRECT—From a local attachment • MSDP—From the Multicast Source Discovery Protocol (MSDP) • DR—From the designated router 	extensive
RP selected	For sparse mode and bidirectional mode, the identity of the RP for the specified group address.	<i>group-address</i>

Sample Output

show pim rps

```

user@host> show pim rps
Instance: PIM.master
Address family INET
RP address      Type      Mode   Holdtime Timeout Groups  Group prefixes
10.10.1.3       static   bidir   150     None     2  224.1.3.0/24
                225.1.3.0/24
10.10.13.2      static   bidir   150     None     2  224.1.1.0/24
                225.1.1.0/24

```

show pim rps brief

The output for the **show pim rps brief** command is identical to that for the **show pim rps** command. For sample output, see [show pim rps on page 548](#).

show pim rps <group-address> (Bidirectional PIM)

```

user@host> show pim rps 224.1.1.1
Instance: PIM.master

224.1.0.0/16
  11.4.12.75 (Bidirectional)

RP selected: 11.4.12.75

```

show pim rps <group-address> (PIM Dense Mode)

```

user@host> show pim rps 224.1.1.1
Instance: PIM.master

Dense Mode active for group 224.1.1.1

```

show pim rps <group-address> (SSM Range Without asm-override-ssm Configured)

```

user@host> show pim rps 224.1.1.1

```

Instance: PIM.master

Source-specific Mode (SSM) active for group 224.1.1.1

show pim rps <group-address> (SSM Range With asm-override-ssm Configured and a Sparse-Mode RP)

user@host> show pim rps 224.1.1.1

Instance: PIM.master

Source-specific Mode (SSM) active with Sparse Mode ASM override for group 224.1.1.1

224.1.0.0/16
11.4.12.75

RP selected: 11.4.12.75

show pim rps <group-address> (SSM Range With asm-override-ssm Configured and a Bidirectional RP)

user@host> show pim rps 224.1.1.1

Instance: PIM.master

Source-specific Mode (SSM) active with Sparse Mode ASM override for group 224.1.1.1

224.1.0.0/16
11.4.12.75 (Bidirectional)

RP selected: (null)

show pim rps instance

user@host> show pim rps instance VPN-A

Instance: PIM.VPN-A

Address family INET

RP address	Type	Holdtime	Timeout	Groups	Group prefixes
10.10.47.100	static	0	None	1	224.0.0.0/4

Address family INET6

show pim rps extensive (PIM Sparse Mode)

user@host> show pim rps extensive

Instance: PIM.master

Family: INET

RP: 10.255.245.91

Learned via: static configuration

Time Active: 00:05:48

Holdtime: 45 with 36 remaining

Device Index: 122

Subunit: 32768

Interface: pd-6/0/0.32768

Group Ranges:

224.0.0.0/4, 36s remaining

Active groups using RP:

225.1.1.1

total 1 groups active

Register State for RP:

Group	Source	FirstHop	RP Address	State	Timeout
225.1.1.1	192.168.195.78	10.255.14.132	10.255.245.91	Receive	0

show pim rps extensive (Bidirectional PIM)

```
user@host> show pim rps extensive
Instance: PIM.master
Address family INET

RP: 10.10.1.3
Learned via: static configuration
Mode: Bidirectional
Time Active: 01:58:07
Holdtime: 150
Group Ranges:
    224.1.3.0/24
    225.1.3.0/24

RP: 10.10.13.2
Learned via: static configuration
Mode: Bidirectional
Time Active: 01:58:07
Holdtime: 150
Group Ranges:
    224.1.1.0/24
    225.1.1.0/24
```

show pim rps extensive (PIM Anycast RP in Use)

```
user@host> show pim rps extensive
Instance: PIM.master

Family: INET
RP: 10.10.10.2
Learned via: static configuration
Time Active: 00:54:52
Holdtime: 0
Device Index: 130
Subunit: 32769
Interface: pimd.32769
Group Ranges:
    224.0.0.0/4
Active groups using RP:
    224.10.10.10

    total 1 groups active

Anycast-PIM rpset:
    10.100.111.34
    10.100.111.17
    10.100.111.55

Anycast-PIM local address used: 10.100.111.1
Anycast-PIM Register State:

```

Group	Source	Origin
224.1.1.1	10.10.95.2	DIRECT
224.1.1.2	10.10.95.2	DIRECT
224.10.10.10	10.10.70.1	MSDP
224.10.10.11	10.10.70.1	MSDP
224.20.20.1	10.10.71.1	DR

```
Address family INET6

Anycast-PIM rpset:
```



```
ab::1
ab::2
Anycast-PIM local address used: cd::1
```

Anycast-PIM Register State:

Group	Source	Origin
::224.1.1.1	::10.10.95.2	DIRECT
::224.1.1.2	::10.10.95.2	DIRECT
::224.20.20.1	::10.10.71.1	DR

show pim snooping interfaces

Syntax	show pim snooping interfaces <brief detail> <instance <i>instance-name</i> > <interface <i>interface-name</i> > <vlan-id <i>vlan-identifier</i> >
Release Information	Command introduced in Junos OS Release 12.3 for MX Series 3D Universal Edge devices. Command introduced in Junos OS Release 13.2 for M Series Multiservice Edge devices.
Description	Display information about PIM snooping interfaces.
Options	<p>none—Display detailed information.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>instance <instance-name>—(Optional) Display PIM snooping interface information for the specified routing instance.</p> <p>interface <interface-name>—(Optional) Display PIM snooping information for the specified interface only.</p> <p>vlan-id <vlan-identifier>—(Optional) Display PIM snooping interface information for the specified VLAN.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> <i>PIM Snooping for VPLS</i>
List of Sample Output	show pim snooping interfaces on page 553 show pim snooping interfaces instance vpls1 on page 553 show pim snooping interfaces interface <interface-name> on page 554 show pim snooping interfaces vlan-id <vlan-id> on page 554
Output Fields	Table 73 on page 552 lists the output fields for the show pim snooping interface command. Output fields are listed in the approximate order in which they appear.

Table 73: show pim snooping interface Output Fields

Field Name	Field Description	Level of Output
Instance	Routing instance for PIM snooping.	All levels
Learning-Domain	Learning domain for snooping.	All levels
Name	Router interfaces that are part of this learning domain.	All levels
State	State of the interface: Up , or Down .	All levels

Table 73: show pim snooping interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
IP-Version	Version of IP used: 4 for IPv4, or 6 for IPv6.	All levels
NbrCnt	Number of neighboring routers connected through the specified interface.	All levels
DR address	IP address of the designated router.	All levels

Sample Output

show pim snooping interfaces

```

user@host> show pim snooping interfaces
Instance: vpls1
Learning-Domain: vlan-id 10
Name State IP-Version NbrCnt
ge-1/3/1.10 Up 4 1
ge-1/3/3.10 Up 4 1
ge-1/3/5.10 Up 4 1
ge-1/3/7.10 Up 4 1
DR address: 20.0.110.5
DR flooding is ON

Learning-Domain: vlan-id 20
Name State IP-Version NbrCnt
ge-1/3/1.20 Up 4 1
ge-1/3/3.20 Up 4 1
ge-1/3/5.20 Up 4 1
ge-1/3/7.20 Up 4 1
DR address: 20.0.120.5
DR flooding is ON

```

show pim snooping interfaces instance vpls1

```

user@host> show pim snooping interfaces instance vpls1
Instance: vpls1

Learning-Domain: vlan-id 10
Name State IP-Version NbrCnt
ge-1/3/1.10 Up 4 1
ge-1/3/3.10 Up 4 1
ge-1/3/5.10 Up 4 1
ge-1/3/7.10 Up 4 1
DR address: 20.0.110.5
DR flooding is ON

Learning-Domain: vlan-id 20
Name State IP-Version NbrCnt
ge-1/3/1.20 Up 4 1
ge-1/3/3.20 Up 4 1
ge-1/3/5.20 Up 4 1
ge-1/3/7.20 Up 4 1
DR address: 20.0.120.5
DR flooding is ON

```

show pim snooping interfaces interface <interface-name>

```
user@host> show pim snooping interfaces interface ge-1/3/1.10
Instance: vpls1
Learning-Domain: vlan-id 10

Name State IP-Version NbrCnt
ge-1/3/1.10 Up 4 1
DR address: 20.0.110.5
DR flooding is ON

Learning-Domain: vlan-id 20
DR address: 20.0.120.5
DR flooding is ON
```

show pim snooping interfaces vlan-id <vlan-id>

```
user@host> show pim snooping interfaces vlan-id 10
Instance: vpls1
Learning-Domain: vlan-id 10

Name State IP-Version NbrCnt
ge-1/3/1.10 Up 4 1
ge-1/3/3.10 Up 4 1
ge-1/3/5.10 Up 4 1
ge-1/3/7.10 Up 4 1
DR address: 20.0.110.5
DR flooding is ON
```

show pim snooping join

Syntax	show pim snooping join <brief detail extensive> <instance <i>instance-name</i> > <vlan-id <i>vlan-id</i> >
Release Information	Command introduced in Junos OS Release 12.3 for MX Series 3D Universal Edge devices. Command introduced in Junos OS Release 13.2 for M Series Multiservice Edge devices.
Description	Display information about Protocol Independent Multicast (PIM) snooping joins.
Options	<p>none—Display detailed information.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display PIM snooping join information for the specified routing instance.</p> <p>vlan-id <i>vlan-identifier</i>—(Optional) Display PIM snooping join information for the specified VLAN.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>PIM Snooping for VPLS</i>
List of Sample Output	show pim snooping join on page 556 show pim snooping join extensive on page 557 show pim snooping join instance on page 557 show pim snooping join vlan-id on page 557
Output Fields	Table 74 on page 555 lists the output fields for the show pim snooping join command. Output fields are listed in the approximate order in which they appear.

Table 74: show pim snooping join Output Fields

Field Name	Field Description	Level of Output
Instance	Routing instance for PIM snooping.	All levels
Learning-Domain	Learning domain for PIM snooping.	All levels
Group	Multicast group address.	All levels
Source	Multicast source address: <ul style="list-style-type: none"> • * (wildcard value) • <ipv4-address> • <ipv6-address> 	All levels

Table 74: show pim snooping join Output Fields (*continued*)

Field Name	Field Description	Level of Output
Flags	PIM flags: <ul style="list-style-type: none"> • bidirectional—Bidirectional mode entry. • dense—Dense mode entry. • rptree—Entry is on the rendezvous point tree. • sparse—Sparse mode entry. • spt—Entry is on the shortest-path tree for the source. • wildcard—Entry is on the shared tree. 	All levels
Upstream state	Information about the upstream interface: <ul style="list-style-type: none"> • Join to RP—Sending a join to the rendezvous point. • Join to Source—Sending a join to the source. • Local RP—Sending neither join messages nor prune messages toward the RP, because this router is the rendezvous point. • Local Source—Sending neither join messages nor prune messages toward the source, because the source is locally attached to this routing device. • Prune to RP—Sending a prune to the rendezvous point. • Prune to Source—Sending a prune to the source. <p>NOTE: RP group range entries have None in the Upstream state field because RP group ranges do not trigger actual PIM join messages between routers.</p>	All levels
Upstream neighbor	Information about the upstream neighbor: Direct , Local , Unknown , or a specific IP address. For bidirectional PIM, Direct means that the interface is directly connected to a subnet that contains a phantom RP address.	All levels
Upstream port	RPF interface toward the source address for the source-specific state (S,G) or toward the rendezvous point (RP) address for the non-source-specific state (*G). For bidirectional PIM, RP Link means that the interface is directly connected to a subnet that contains a phantom RP address.	All levels
Downstream port	Information about downstream interfaces.	extensive
Downstream neighbors	Address of the downstream neighbor.	extensive
Timeout	Time remaining until the downstream join state is updated (in seconds).	extensive

Sample Output

show pim snooping join

```

user@host> show pim snooping join
Instance: vpls1

Learning-Domain: vlan-id 10

```

```

Group: 225.1.1.2
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 20.0.110.4, port: ge-1/3/5.10

```

```

Learning-Domain: vlan-id 20
Group: 225.1.1.3
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 20.0.120.4, port: ge-1/3/5.20

```

show pim snooping join extensive

```

user@host> show pim snooping join extensive
Instance: vpls1
Learning-Domain: vlan-id 10

Group: 225.1.1.2
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 20.0.110.4, port: ge-1/3/5.10
Downstream port: ge-1/3/1.10
Downstream neighbors:
20.0.110.2 State: Join Flags: SRW Timeout: 166

Learning-Domain: vlan-id 20
Group: 225.1.1.3
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 20.0.120.4, port: ge-1/3/5.20
Downstream port: ge-1/3/3.20
Downstream neighbors:
20.0.120.3 State: Join Flags: SRW Timeout: 168

```

show pim snooping join instance

```

user@host> show pim snooping join instance vpls1
Instance: vpls1

Learning-Domain: vlan-id 10
Group: 225.1.1.2
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 20.0.110.4, port: ge-1/3/5.10

Learning-Domain: vlan-id 20
Group: 225.1.1.3
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 20.0.120.4, port: ge-1/3/5.20

```

show pim snooping join vlan-id

```

user@host> show pim snooping join vlan-id 10

```

Instance: vpls1
Learning-Domain: vlan-id 10
Group: 225.1.1.2
Source: *
Flags: sparse,rptree,wildcard
Upstream state: None
Upstream neighbor: 20.0.110.4, port: ge-1/3/5.10

show pim snooping neighbors

Syntax	show pim snooping neighbors <brief detail> <instance <i>instance-name</i> > <interface <i>interface-name</i> > <vlan-id <i>vlan-identifier</i> >
Release Information	Command introduced in Junos OS Release 12.3 for MX Series 3D Universal Edge devices. Command introduced in Junos OS Release 13.2 for M Series Multiservice Edge devices.
Description	Display information about Protocol Independent Multicast (PIM) snooping neighbors.
Options	<p>none—Display detailed information.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display PIM snooping neighbor information for the specified routing instance.</p> <p>interface <i>interface-name</i>—(Optional) Display information for the specified PIM snooping neighbor interface.</p> <p>vlan-id <i>vlan-identifier</i>—(Optional) Display PIM snooping neighbor information for the specified VLAN.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Interface Priority for PIM Designated Router Selection</i> • <i>Modifying the PIM Hello Interval</i> • <i>PIM Snooping for VPLS</i> • show pim neighbors on page 541
List of Sample Output	show pim snooping neighbors on page 560 show pim snooping neighbors detail on page 561 show pim snooping neighbors instance on page 562 show pim snooping neighbors interface on page 562 show pim snooping neighbors vlan-id on page 563
Output Fields	<p>Table 75 on page 559 lists the output fields for the show pim snooping neighbors command. Output fields are listed in the approximate order in which they appear.</p>

Table 75: show pim snooping neighbors Output Fields

Field Name	Field Description	Level of Output
Instance	Routing instance for PIM snooping.	All levels
Learning-Domain	Learning domain for PIM snooping.	All levels

Table 75: show pim snooping neighbors Output Fields (*continued*)

Field Name	Field Description	Level of Output
Interface	Router interface for which PIM snooping neighbor details are displayed.	All levels
Option	PIM snooping options available on the specified interface: <ul style="list-style-type: none"> • H = Hello Option Holdtime • P = Hello Option DR Priority • L = Hello Option LAN Prune Delay • G = Generation Identifier • T = Tracking Bit 	All levels
Uptime	Time the neighbor has been operational since the PIM process was last initialized, in the format dd:hh:mm:ss ago for less than a week and nwnd:hh:mm:ss ago for more than a week.	All levels
Neighbor addr	IP address of the PIM snooping neighbor connected through the specified interface.	All levels
Address	IP address of the specified router interface.	All levels
Hello Option Holdtime	Time for which the neighbor is available, in seconds. The range of values is 0 through 65,535 .	detail
Hello Option DR Priority	Designated router election priority. The range of values is 0 through 255 . NOTE: By default, every PIM interface has the lowest probability (priority 0) of being selected as the DR.	detail
Hello Option Generation ID	9-digit or 10-digit number used to tag hello messages.	detail
Hello Option LAN Prune Delay	Time to wait before the neighbor receives prune messages, in the format delay nnn ms override nnnn ms .	detail

Sample Output

show pim snooping neighbors

```
user@host> show pim snooping neighbors
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking Bit
```

```
Instance: vpls1
Learning-Domain: vlan-id 10
```

```
Interface Option Uptime Neighbor addr
ge-1/3/1.10 HPLGT 00:43:33 20.0.110.2
ge-1/3/3.10 HPLGT 00:43:33 20.0.110.3
ge-1/3/5.10 HPLGT 00:43:33 20.0.110.4
ge-1/3/7.10 HPLGT 00:43:33 20.0.110.5
```

```
Learning-Domain: vlan-id 20
```

```

Interface Option Uptime Neighbor addr
ge-1/3/1.20 HPLGT 00:43:33 20.0.120.2
ge-1/3/3.20 HPLGT 00:43:33 20.0.120.3
ge-1/3/5.20 HPLGT 00:43:33 20.0.120.4
ge-1/3/7.20 HPLGT 00:43:33 20.0.120.5

```

show pim snooping neighbors detail

```

user@host> show pim snooping neighbors detail
Instance: vpls1
Learning-Domain: vlan-id 10

Interface: ge-1/3/1.10
Address: 20.0.110.2
Uptime: 00:44:51
Hello Option Holdtime: 105 seconds 83 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 830908833
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Tracking is supported

Interface: ge-1/3/3.10
Address: 20.0.110.3
Uptime: 00:44:51
Hello Option Holdtime: 105 seconds 97 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 2056520742
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Tracking is supported

Interface: ge-1/3/5.10
Address: 20.0.110.4
Uptime: 00:44:51
Hello Option Holdtime: 105 seconds 81 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 1152066227
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Tracking is supported

Interface: ge-1/3/7.10
Address: 20.0.110.5
Uptime: 00:44:51
Hello Option Holdtime: 105 seconds 96 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 1113200338
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Tracking is supported
Learning-Domain: vlan-id 20

Interface: ge-1/3/1.20
Address: 20.0.120.2
Uptime: 00:44:51
Hello Option Holdtime: 105 seconds 81 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 963205167
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Tracking is supported

Interface: ge-1/3/3.20
Address: 20.0.120.3

```

```
Uptime: 00:44:51
Hello Option Holdtime: 105 seconds 104 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 166921538
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Tracking is supported
```

```
Interface: ge-1/3/5.20
Address: 20.0.120.4
Uptime: 00:44:51
Hello Option Holdtime: 105 seconds 88 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 789422835
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Tracking is supported
```

```
Interface: ge-1/3/7.20
Address: 20.0.120.5
Uptime: 00:44:51
Hello Option Holdtime: 105 seconds 88 remaining
Hello Option DR Priority: 1
Hello Option Generation ID: 1563649680
Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
Tracking is supported
```

show pim snooping neighbors instance

```
user@host> show pim snooping neighbors instance vpls1
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking Bit
```

```
Instance: vpls1
Learning-Domain: vlan-id 10
```

```
Interface Option Uptime Neighbor addr
ge-1/3/1.10 HPLGT 00:46:03 20.0.110.2
ge-1/3/3.10 HPLGT 00:46:03 20.0.110.3
ge-1/3/5.10 HPLGT 00:46:03 20.0.110.4
ge-1/3/7.10 HPLGT 00:46:03 20.0.110.5
```

```
Learning-Domain: vlan-id 20
```

```
Interface Option Uptime Neighbor addr
ge-1/3/1.20 HPLGT 00:46:03 20.0.120.2
ge-1/3/3.20 HPLGT 00:46:03 20.0.120.3
ge-1/3/5.20 HPLGT 00:46:03 20.0.120.4
ge-1/3/7.20 HPLGT 00:46:03 20.0.120.5
```

show pim snooping neighbors interface

```
user@host> show pim snooping neighbors interface ge-1/3/1.20
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking Bit
```

```
Instance: vpls1
Learning-Domain: vlan-id 10
Learning-Domain: vlan-id 20
```

```
Interface Option Uptime Neighbor addr
ge-1/3/1.20 HPLGT 00:48:04 20.0.120.2
```

show pim snooping neighbors vlan-id

```
user@host> show pim snooping neighbors vlan-id 10
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority, T = Tracking Bit
```

```
Instance: vpls1
Learning-Domain: vlan-id 10
```

```
Interface Option Uptime Neighbor addr
ge-1/3/1.10 HPLGT 00:49:12 20.0.110.2
ge-1/3/3.10 HPLGT 00:49:12 20.0.110.3
ge-1/3/5.10 HPLGT 00:49:12 20.0.110.4
ge-1/3/7.10 HPLGT 00:49:12 20.0.110.5
```

show pim snooping statistics

Syntax	show pim snooping statistics <instance <i>instance-name</i> > <interface <i>interface-name</i> > <vlan-id <i>vlan-id</i> >
Release Information	Command introduced in Junos OS Release 12.3 for MX Series 3D Universal Edge devices. Command introduced in Junos OS Release 13.2 for M Series Multiservice Edge devices.
Description	Display Protocol Independent Multicast (PIM) snooping statistics.
Options	<p>none—Display PIM statistics.</p> <p>instance <i>instance-name</i>—(Optional) Display statistics for a specific routing instance enabled by Protocol Independent Multicast (PIM) snooping.</p> <p>interface <i>interface-name</i>—(Optional) Display statistics about the specified interface for PIM snooping.</p> <p>vlan-id <i>vlan-identifier</i>—(Optional) Display PIM snooping statistics information for the specified VLAN.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>PIM Snooping for VPLS</i> • clear pim snooping statistics on page 393
List of Sample Output	show pim snooping statistics on page 565 show pim snooping statistics instance on page 566 show pim snooping statistics interface on page 567 show pim snooping statistics vlan-id on page 567
Output Fields	Table 76 on page 564 lists the output fields for the show pim snooping statistics command. Output fields are listed in the approximate order in which they appear.

Table 76: show pim snooping statistics Output Fields

Field Name	Field Description	Level of Output
Instance	Routing instance for PIM snooping.	All levels
Learning-Domain	Learning domain for PIM snooping.	All levels
Tx J/P messages	Total number of transmitted join/prune packets.	All levels
RX J/P messages	Total number of received join/prune packets.	All levels

Table 76: show pim snooping statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
Rx J/P messages -- seen	Number of join/prune packets seen but not received on the upstream interface.	All levels
Rx J/P messages -- received	Number of join/prune packets received on the downstream interface.	All levels
Rx Hello messages	Total number of received hello packets.	All levels
Rx Version Unknown	Number of packets received with an unknown version number.	All levels
Rx Neighbor Unknown	Number of packets received from an unknown neighbor.	All levels
Rx Upstream Neighbor Unknown	Number of packets received with unknown upstream neighbor information.	All levels
Rx Bad Length	Number of packets received containing incorrect length information.	All levels
Rx J/P Busy Drop	Number of join/prune packets dropped while the router is busy.	All levels
Rx J/P Group Aggregate 0	Number of join/prune packets received containing the aggregate group information.	All levels
Rx Malformed Packet	Number of malformed packets received.	All levels
Rx No PIM Interface	Number of packets received without the interface information.	All levels
Rx No Upstream Neighbor	Number of packets received without upstream neighbor information.	All levels
Rx Unknown Hello Option	Number of hello packets received with unknown options.	All levels

Sample Output

show pim snooping statistics

```

user@host> show pim snooping statistics
Instance: vpls1
Learning-Domain: vlan-id 10

Tx J/P messages 0
RX J/P messages 8
Rx J/P messages -- seen 0
Rx J/P messages -- received 8
Rx Hello messages 37
Rx Version Unknown 0
Rx Neighbor Unknown 0

```

```
Rx Upstream Neighbor Unknown 0
Rx Bad Length 0
Rx J/P Busy Drop 0
Rx J/P Group Aggregate 0
Rx Malformed Packet 0
Rx No PIM Interface 0
Rx No Upstream Neighbor 0
Rx Bad Length 0
Rx Neighbor Unknown 0
Rx Unknown Hello Option 0
Rx Malformed Packet 0
```

Learning-Domain: vlan-id 20

```
Tx J/P messages 0
RX J/P messages 2
Rx J/P messages -- seen 0
Rx J/P messages -- received 2
Rx Hello messages 39
Rx Version Unknown 0
Rx Neighbor Unknown 0
Rx Upstream Neighbor Unknown 0
Rx Bad Length 0
Rx J/P Busy Drop 0
Rx J/P Group Aggregate 0
Rx Malformed Packet 0
Rx No PIM Interface 0
Rx No Upstream Neighbor 0
Rx Bad Length 0
Rx Neighbor Unknown 0
Rx Unknown Hello Option 0
Rx Malformed Packet 0
```

show pim snooping statistics instance

```
user@host> show pim snooping statistics instance vpls1
```

Instance: vpls1

Learning-Domain: vlan-id 10

```
Tx J/P messages 0
RX J/P messages 9
Rx J/P messages -- seen 0
Rx J/P messages -- received 9
Rx Hello messages 45
Rx Version Unknown 0
Rx Neighbor Unknown 0
Rx Upstream Neighbor Unknown 0
Rx Bad Length 0
Rx J/P Busy Drop 0
Rx J/P Group Aggregate 0
Rx Malformed Packet 0
Rx No PIM Interface 0
Rx No Upstream Neighbor 0
Rx Bad Length 0
Rx Neighbor Unknown 0
Rx Unknown Hello Option 0
Rx Malformed Packet 0
```

Learning-Domain: vlan-id 20


```

Tx J/P messages 0
RX J/P messages 3
Rx J/P messages -- seen 0
Rx J/P messages -- received 3
Rx Hello messages 47
Rx Version Unknown 0
Rx Neighbor Unknown 0
Rx Upstream Neighbor Unknown 0
Rx Bad Length 0
Rx J/P Busy Drop 0
Rx J/P Group Aggregate 0
Rx Malformed Packet 0
Rx No PIM Interface 0
Rx No Upstream Neighbor 0
Rx Bad Length 0
Rx Neighbor Unknown 0
Rx Unknown Hello Option 0
Rx Malformed Packet 0

```

show pim snooping statistics interface

```

user@host> show pim snooping statistics interface ge-1/3/1.20
Instance: vpls1
Learning-Domain: vlan-id 10
Learning-Domain: vlan-id 20

PIM Interface statistics for ge-1/3/1.20
Tx J/P messages 0
RX J/P messages 0
Rx J/P messages -- seen 0
Rx J/P messages -- received 0
Rx Hello messages 13
Rx Version Unknown 0
Rx Neighbor Unknown 0
Rx Upstream Neighbor Unknown 0
Rx Bad Length 0
Rx J/P Busy Drop 0
Rx J/P Group Aggregate 0
Rx Malformed Packet 0

```

show pim snooping statistics vlan-id

```

user@host> show pim snooping statistics vlan-id 10
Instance: vpls1
Learning-Domain: vlan-id 10

Tx J/P messages 0
RX J/P messages 11
Rx J/P messages -- seen 0
Rx J/P messages -- received 11
Rx Hello messages 64
Rx Version Unknown 0
Rx Neighbor Unknown 0
Rx Upstream Neighbor Unknown 0
Rx Bad Length 0
Rx J/P Busy Drop 0
Rx J/P Group Aggregate 0
Rx Malformed Packet 0
Rx No PIM Interface 0
Rx No Upstream Neighbor 0

```

Rx Bad Length 0
Rx Neighbor Unknown 0

show pim source

Syntax	<pre>show pim source <brief detail> <inet inet6> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <source-prefix></pre>
Syntax (EX Series Switch and the QFX Series)	<pre>show pim source <brief detail> <inet inet6> <instance <i>instance-name</i>> <source-prefix></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Display information about the Protocol Independent Multicast (PIM) source reverse path forwarding (RPF) state.
Options	<p>none—Display standard information about the PIM RPF state for all supported family addresses for all routing instances.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>inet inet6—(Optional) Display information for IPv4 or IPv6 family addresses, respectively.</p> <p>instance <i>instance-name</i>—(Optional) Display information about the RPF state for a specific PIM-enabled routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>source-prefix—(Optional) Display the state for source RPF states in the given range.</p>
Required Privilege Level	view
List of Sample Output	<p>show pim source on page 570</p> <p>show pim source brief on page 570</p> <p>show pim source detail on page 570</p> <p>show pim source (Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 571</p>
Output Fields	Table 77 on page 570 describes the output fields for the show pim source command. Output fields are listed in the approximate order in which they appear.

Table 77: show pim source Output Fields

Field Name	Field Description
Instance	Name of the routing instance.
Source	Address of the source or reverse path.
Prefix/length	Prefix and prefix length for the route used to reach the RPF address.
Upstream Protocol	Protocol toward the source address.
Upstream interface	RPF interface toward the source address. A pseudo multipoint LDP (M-LDP) interface appears on egress nodes in M-LDP point-to-multipoint LSPs with inband signaling.
Upstream Neighbor	Address of the RPF neighbor used to reach the source address. The multipoint LDP (M-LDP) root appears on egress nodes in M-LDP point-to-multipoint LSPs with inband signaling.

Sample Output

show pim source

```

user@host> show pim source
Instance: PIM.master Family: INET

Source 10.255.14.144
  Prefix 10.255.14.144/32
  Upstream interface Local
  Upstream neighbor Local

Source 10.255.70.15
  Prefix 10.255.70.15/32
  Upstream interface so-1/0/0.0
  Upstream neighbor 10.111.10.2

Instance: PIM.master Family: INET6

```

show pim source brief

The output for the **show pim source brief** command is identical to that for the **show pim source** command. For sample output, see [show pim source on page 570](#).

show pim source detail

```

user@host> show pim source detail
Instance: PIM.master Family: INET

Source 10.255.14.144
  Prefix 10.255.14.144/32
  Upstream interface Local
  Upstream neighbor Local
  Active groups:228.0.0.0
  239.1.1.1

```

239.1.1.1

Source 10.255.70.15
 Prefix 10.255.70.15/32
 Upstream interface so-1/0/0.0
 Upstream neighbor 10.111.10.2
 Active groups:239.1.1.1

Instance: PIM.master Family: INET6

show pim source (Egress Node with Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

user@host> show pim source

Instance: PIM.master Family: INET

Source 1.1.1.1
 Prefix 1.1.1.1/32
 Upstream interface Local
 Upstream neighbor Local

Source 1.2.7.7
 Prefix 1.2.7.0/24
 Upstream protocol MLDP
 Upstream interface Pseudo MLDP
 Upstream neighbor MLDP LSP root <1.1.1.2>

Source 192.168.219.11
 Prefix 192.168.219.0/28
 Upstream protocol MLDP
 Upstream interface Pseudo MLDP
 Upstream neighbor MLDP LSP root <1.1.1.2>

Instance: PIM.master Family: INET6

Source abcd::1:2:7:7
 Prefix abcd::1:2:7:0/120
 Upstream protocol MLDP
 Upstream interface Pseudo MLDP
 Upstream neighbor MLDP LSP root <1.1.1.2>

show pim statistics

Syntax	show pim statistics <inet inet6> <instance <i>instance-name</i> > <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch and the QFX Series)	show pim statistics <inet inet6> <instance <i>instance-name</i> > <interface <i>interface-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. inet6 and instance options introduced in Junos OS Release 10.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series. Support for bidirectional PIM added in Junos OS Release 12.1.
Description	Display Protocol Independent Multicast (PIM) statistics.
Options	none —Display PIM statistics. inet inet6 —(Optional) Display IPv4 or IPv6 PIM statistics, respectively. instance <i>instance-name</i> —(Optional) Display statistics for a specific routing instance enabled by Protocol Independent Multicast (PIM). interface <i>interface-name</i> —(Optional) Display statistics about the specified interface. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• clear pim statistics on page 395
List of Sample Output	show pim statistics on page 579 show pim statistics inet interface <interface-name> on page 581 show pim statistics inet6 interface <interface-name> on page 581 show pim statistics instance <instance-name> on page 582 show pim statistics interface <interface-name> on page 583
Output Fields	Table 78 on page 573 describes the output fields for the show pim statistics command. Output fields are listed in the approximate order in which they appear.

Table 78: show pim statistics Output Fields

Field Name	Field Description
Instance	<p>Name of the routing instance.</p> <p>This field only appears if you specify an interface, for example:</p> <ul style="list-style-type: none"> • inet interface <i>interface-name</i> • inet6 interface <i>interface-name</i> • interface <i>interface-name</i>
Family	<p>Output is for IPv4 or IPv6 PIM statistics. INET indicates IPv4 statistics, and INET6 indicates IPv6 statistics.</p> <p>This field only appears if you specify an interface, for example:</p> <ul style="list-style-type: none"> • inet interface <i>interface-name</i> • inet6 interface <i>interface-name</i> • interface <i>interface-name</i>
PIM statistics	PIM statistics for all interfaces or for the specified interface.
PIM message type	Message type for which statistics are displayed.
Received	Number of received statistics.
Sent	Number of messages sent of a certain type.
Rx errors	Number of received packets that contained errors.
V2 Hello	PIM version 2 hello packets.
V2 Register	PIM version 2 register packets.
V2 Register Stop	PIM version 2 register stop packets.
V2 Join Prune	PIM version 2 join and prune packets.
V2 Bootstrap	PIM version 2 bootstrap packets.
V2 Assert	PIM version 2 assert packets.
V2 Graft	PIM version 2 graft packets.
V2 Graft Ack	PIM version 2 graft acknowledgment packets.
V2 Candidate RP	PIM version 2 candidate RP packets.
V2 State Refresh	<p>PIM version 2 control messages related to PIM dense mode (PIM-DM) state refresh.</p> <p>State refresh is an extension to PIM-DM. It not supported in Junos OS.</p>

Table 78: show pim statistics Output Fields (*continued*)

Field Name	Field Description
V2 DF Election	PIM version 2 send and receive messages associated with bidirectional PIM designated forwarder election.
V1 Query	PIM version 1 query packets.
V1 Register	PIM version 1 register packets.
V1 Register Stop	PIM version 1 register stop packets.
V1 Join Prune	PIM version 1 join and prune packets.
V1 RP Reachability	PIM version 1 RP reachability packets.
V1 Assert	PIM version 1 assert packets.
V1 Graft	PIM version 1 graft packets.
V1 Graft Ack	PIM version 1 graft acknowledgment packets.
AutoRP Announce	Auto-RP announce packets.
AutoRP Mapping	Auto-RP mapping packets.
AutoRP Unknown type	Auto-RP packets with an unknown type.
Anycast Register	Auto-RP announce packets.
Anycast Register Stop	Auto-RP announce packets.
Global Statistics	Summary of PIM statistics for all interfaces.
Hello dropped on neighbor policy	Number of hello packets dropped because of a configured neighbor policy.
Unknown type	Number of PIM control packets received with an unknown type.
V1 Unknown type	Number of PIM version 1 control packets received with an unknown type.
Unknown Version	Number of PIM control packets received with an unknown version. The version is not version 1 or version 2.
Neighbor unknown	Number of PIM control packets received (excluding PIM hello) without first receiving the hello packet.
Bad Length	Number of PIM control packets received for which the packet size does not match the PIM length field in the packet.

Table 78: show pim statistics Output Fields (*continued*)

Field Name	Field Description
Bad Checksum	Number of PIM control packets received for which the calculated checksum does not match the checksum field in the packet.
Bad Receive If	Number of PIM control packets received on an interface that does not have PIM configured.
Rx Bad Data	Number of PIM control packets received that contain data for TCP Bad register packets.
Rx Intf disabled	Number of PIM control packets received on an interface that has PIM disabled.
Rx V1 Require V2	Number of PIM version 1 control packets received on an interface configured for PIM version 2.
Rx V2 Require V1	Number of PIM version 2 control packets received on an interface configured for PIM version 1.
Rx Register not RP	Number of PIM register packets received when the routing device is not the RP for the group.
Rx Register no route	Number of PIM register packets received when the RP does not have a unicast route back to the source.
Rx Register no decap if	Number of PIM register packets received when the RP does not have a de-encapsulation interface.
Null Register Timeout	Number of NULL register timeout packets.
RP Filtered Source	Number of PIM packets received when the routing device has a source address filter configured for the RP.
Rx Unknown Reg Stop	Number of register stop messages received with an unknown type.
Rx Join/Prune no state	Number of join and prune messages received for which the routing device has no state.
Rx Join/Prune on upstream if	Number of join and prune messages received on the interface used to reach the upstream routing device, toward the RP.
Rx Join/Prune for invalid group	Number of join or prune messages received for invalid multicast group addresses.
Rx Join/Prune messages dropped	Number of join and prune messages received and dropped.
Rx sparse join for dense group	Number of PIM sparse mode join messages received for a group that is configured for dense mode.

Table 78: show pim statistics Output Fields (*continued*)

Field Name	Field Description
Rx Graft/Graft Ack no state	Number of graft and graft acknowledgment messages received for which the router or switch has no state.
Rx Graft on upstream if	Number of graft messages received on the interface used to reach the upstream routing device, toward the RP.
Rx CRP not BSR	Number of BSR messages received in which the PIM message type is Candidate-RP-Advertisement, not Bootstrap.
Rx BSR when BSR	Number of BSR messages received in which the PIM message type is Bootstrap.
Rx BSR not RPF if	Number of BSR messages received on an interface that is not the RPF interface.
Rx unknown hello opt	Number of PIM hello packets received with options that Junos OS does not support.
Rx data no state	Number of PIM control packets received for which the routing device has no state for the data type.
Rx RP no state	Number of PIM control packets received for which the routing device has no state for the RP.
Rx aggregate	Number of PIM aggregate MDT packets received.
Rx malformed packet	Number of PIM control packets received with a malformed IP unicast or multicast address family.
No RP	Number of PIM control packets received with no RP address.
No register encaps if	Number of PIM register packets received when the first-hop routing device does not have an encapsulation interface.
No route upstream	Number of PIM control packets received when the routing device does not have a unicast route to the the interface used to reach the upstream routing device, toward the RP.
Nexthop Unusable	Number of PIM control packets with an unusable nexthop. A path can be unusable if the route is hidden or the link is down.
RP mismatch	Number of PIM control packets received for which the routing device has an RP mismatch.
RP mode mismatch	RP mode (sparse or bidirectional) mismatches encountered when processing join and prune messages.
RPF neighbor unknown	Number of PIM control packets received for which the routing device has an unknown RPF neighbor for the source.

Table 78: show pim statistics Output Fields (*continued*)

Field Name	Field Description
Rx Joins/Prunes filtered	The number of join and prune messages filtered because of configured route filters and source address filters.
Tx Joins/Prunes filtered	The number of join and prune messages filtered because of configured route filters and source address filters.
Embedded-RP invalid addr	Number of packets received with an invalid embedded RP address in PIM join messages and other types of messages sent between routing domains.
Embedded-RP limit exceed	Number of times the limit configured with the maximum-rps statement is exceeded. The maximum-rps statement limits the number of embedded RPs created in a specific routing instance. The range is from 1 through 500. The default is 100.
Embedded-RP added	<p>Number of packets in which the embedded RP for IPv6 is added.</p> <p>The following receive events trigger extraction of an IPv6 embedded RP address on the routing device:</p> <ul style="list-style-type: none"> • Multicast Listener Discovery (MLD) report for an embedded RP multicast group address • PIM join message with an embedded RP multicast group address • Static embedded RP multicast group address associated with an interface • Packets sent to an embedded RP multicast group address received on the DR <p>An embedded RP node discovered through these receive events is added if it does not already exist on the routing platform.</p>
Embedded-RP removed	Number of packets in which the embedded RP for IPv6 is removed. The embedded RP is removed whenever all PIM join states using this RP are removed or the configuration changes to remove the embedded RP feature.
Rx Register msgs filtering drop	Number of received register messages dropped because of a filter configured for PIM register messages.
Tx Register msgs filtering drop	Number of register messages dropped because of a filter configured for PIM register messages.
Rx Bidir Join/Prune on non-Bidir if	Error counter for join and prune messages received on non-bidirectional PIM interfaces.
Rx Bidir Join/Prune on non-DF if	Error counter for join and prune messages received on non-designated forwarder interfaces.
V4 (S,G) Maximum	Maximum number of (S,G) IPv4 multicast routes accepted for the VPN routing and forwarding (VRF) routing instance. If this number is met, additional (S,G) entries are not accepted.

Table 78: show pim statistics Output Fields (*continued*)

Field Name	Field Description
V4 (S,G) Accepted	Number of accepted (S,G) IPv4 multicast routes.
V4 (S,G) Threshold	Threshold at which a warning message is logged (percentage of the maximum number of (S,G) IPv4 multicast routes accepted by the device).
V4 (S,G) Log Interval	Time (in seconds) between consecutive log messages.
V6 (S,G) Maximum	Maximum number of (S,G) IPv6 multicast routes accepted for the VPN routing and forwarding (VRF) routing instance. If this number is met, additional (S,G) entries are not accepted.
V6 (S,G) Accepted	Number of accepted (S,G) IPv6 multicast routes.
V6 (S,G) Threshold	Threshold at which a warning message is logged (percentage of the maximum number of (S,G) IPv6 multicast routes accepted by the device).
V6 (S,G) Log Interval	Time (in seconds) between consecutive log messages.
V4 (grp-prefix, RP) Maximum	Maximum number of group-to-rendezvous point (RP) IPv4 multicast mappings accepted for the VRF routing instance. If this number is met, additional mappings are not accepted.
V4 (grp-prefix, RP) Accepted	Number of accepted group-to-RP IPv4 multicast mappings.
V4 (grp-prefix, RP) Threshold	Threshold at which a warning message is logged (percentage of the maximum number of group-to-RP IPv4 multicast mappings accepted by the device).
V4 (grp-prefix, RP) Log Interval	Time (in seconds) between consecutive log messages.
V6 (grp-prefix, RP) Maximum	Maximum number of group-to RP IPv6 multicast mappings accepted for the VRF routing instance. If this number is met, additional mappings are not accepted.
V6 (grp-prefix, RP) Accepted	Number of accepted group-to-RP IPv6 multicast mappings.
V6 (grp-prefix, RP) Threshold	Threshold at which a warning message is logged (percentage of the maximum number of group-to-RP IPv6 multicast mappings accepted by the device).
V6 (grp-prefix, RP) Log Interval	Time (in seconds) between consecutive log messages.

Table 78: show pim statistics Output Fields (*continued*)

Field Name	Field Description
V4 Register Maximum	Maximum number of IPv4 PIM registers accepted for the VRF routing instance. If this number is met, additional PIM registers are not accepted. You configure the register limits on the RP.
V4 Register Accepted	Number of accepted IPv4 PIM registers.
V4 Register Threshold	Threshold at which a warning message is logged (percentage of the maximum number of IPv4 PIM registers accepted by the device).
V4 Register Log Interval	Time (in seconds) between consecutive log messages.
V6 Register Maximum	Maximum number of IPv6 PIM registers accepted for the VRF routing instance. If this number is met, additional PIM registers are not accepted. You configure the register limits on the RP.
V6 Register Accepted	Number of accepted IPv6 PIM registers.
V6 Register Threshold	Threshold at which a warning message is logged (percentage of the maximum number of IPv6 PIM registers accepted by the device).
V6 Register Log Interval	Time (in seconds) between consecutive log messages.
(*G) Join drop due to SSM range check	PIM join messages that are dropped because the multicast addresses are outside of the SSM address range of 232.0.0.0 through 232.255.255.255. You can extend the accepted SSM address range by configuring the ssm-groups statement.

Sample Output

show pim statistics

```

user@host> show pim statistics
PIM Message type    Received    Sent    Rx errors
V2 Hello            15          32         0
V2 Register          0          362        0
V2 Register Stop     483         0         0
V2 Join Prune        18          518        0
V2 Bootstrap         0           0         0
V2 Assert            0           0         0
V2 Graft             0           0         0
V2 Graft Ack         0           0         0
V2 Candidate RP      0           0         0
V2 State Refresh     0           0         0
V2 DF Election       0           0         0
V1 Query             0           0         0
V1 Register          0           0         0
V1 Register Stop     0           0         0
V1 Join Prune        0           0         0

```

V1 RP Reachability	0	0	0
V1 Assert	0	0	0
V1 Graft	0	0	0
V1 Graft Ack	0	0	0
AutoRP Announce	0	0	0
AutoRP Mapping	0	0	0
AutoRP Unknown type	0		
Anycast Register	0	0	0
Anycast Register Stop	0	0	0

Global Statistics

Hello dropped on neighbor policy	0
Unknown type	0
V1 Unknown type	0
Unknown Version	0
ipv4 BSR pkt drop due to excessive rate	0
ipv6 BSR pkt drop due to excessive rate	0
Neighbor unknown	0
Bad Length	0
Bad Checksum	0
Bad Receive If	0
Rx Bad Data	0
Rx Intf disabled	0
Rx V1 Require V2	0
Rx V2 Require V1	0
Rx Register not RP	0
Rx Register no route	0
Rx Register no decap if	0
Null Register Timeout	0
RP Filtered Source	0
Rx Unknown Reg Stop	0
Rx Join/Prune no state	0
Rx Join/Prune on upstream if	0
Rx Join/Prune for invalid group	5
Rx Join/Prune messages dropped	0
Rx sparse join for dense group	0
Rx Graft/Graft Ack no state	0
Rx Graft on upstream if	0
Rx CRP not BSR	0
Rx BSR when BSR	0
Rx BSR not RPF if	0
Rx unknown hello opt	0
Rx data no state	0
Rx RP no state	0
Rx aggregate	0
Rx malformed packet	0
Rx illegal TTL	0
Rx illegal destination address	0
No RP	0
No register encap if	0
No route upstream	0
Nexthop Unusable	0
RP mismatch	0
RP mode mismatch	0
RPF neighbor unknown	0
Rx Joins/Prunes filtered	0
Tx Joins/Prunes filtered	0
Embedded-RP invalid addr	0
Embedded-RP limit exceed	0
Embedded-RP added	0

```

Embedded-RP removed                0
Rx Register msgs filtering drop      0
Tx Register msgs filtering drop      0
Rx Bidir Join/Prune on non-Bidir if  0
Rx Bidir Join/Prune on non-DF if     0
(*,G) Join drop due to SSM range check 0

```

Sample Output

show pim statistics inet interface <interface-name>

```

user@host> show pim statistics inet interface ge-0/3/0.0
Instance: PIM.master Family: INET

```

PIM Interface statistics for ge-0/3/0.0

PIM Message type	Received	Sent	Rx errors
V2 Hello	0	4	0
V2 Register	0	0	0
V2 Register Stop	0	0	0
V2 Join Prune	0	0	0
V2 Bootstrap	0	0	0
V2 Assert	0	0	0
V2 Graft	0	0	0
V2 Graft Ack	0	0	0
V2 Candidate RP	0	0	0
V1 Query	0	0	0
V1 Register	0	0	0
V1 Register Stop	0	0	0
V1 Join Prune	0	0	0
V1 RP Reachability	0	0	0
V1 Assert	0	0	0
V1 Graft	0	0	0
V1 Graft Ack	0	0	0
AutoRP Announce	0	0	0
AutoRP Mapping	0	0	0
AutoRP Unknown type	0		
Anycast Register	0	0	0
Anycast Register Stop	0	0	0

Sample Output

show pim statistics inet6 interface <interface-name>

```

user@host> show pim statistics inet6 interface ge-0/3/0.0
Instance: PIM.master Family: INET6

```

PIM Interface statistics for ge-0/3/0.0

PIM Message type	Received	Sent	Rx errors
V2 Hello	0	4	0
V2 Register	0	0	0
V2 Register Stop	0	0	0
V2 Join Prune	0	0	0
V2 Bootstrap	0	0	0
V2 Assert	0	0	0
V2 Graft	0	0	0
V2 Graft Ack	0	0	0
V2 Candidate RP	0	0	0
Anycast Register	0	0	0
Anycast Register Stop	0	0	0

show pim statistics instance <instance-name>

```

user@host> show pim statistics instance VPN-A
PIM Message type      Received      Sent  Rx errors
V2 Hello               31           37      0
V2 Register            0            0      0
V2 Register Stop       0            0      0
V2 Join Prune          0           16      0
V2 Bootstrap           0            0      0
V2 Assert              0            0      0
V2 Graft               0            0      0
V2 Graft Ack           0            0      0
V2 Candidate RP        0            0      0
V2 State Refresh       0            0      0
V2 DF Election         0            0      0
V1 Query               0            0      0
V1 Register            0            0      0
V1 Register Stop       0            0      0
V1 Join Prune          0            0      0
V1 RP Reachability     0            0      0
V1 Assert              0            0      0
V1 Graft               0            0      0
V1 Graft Ack           0            0      0
AutoRP Announce        0            0      0
AutoRP Mapping          0            0      0
AutoRP Unknown type    0            0      0
Anycast Register       0            0      0
Anycast Register Stop  0            0      0

```

Global Statistics

```

Hello dropped on neighbor policy      0
Unknown type                          0
V1 Unknown type                       0
Unknown Version                       0
Neighbor unknown                      0
Bad Length                            0
Bad Checksum                          0
Bad Receive If                        0
Rx Bad Data                           0
Rx Intf disabled                       0
Rx V1 Require V2                      0
Rx V2 Require V1                      0
Rx Register not RP                    0
Rx Register no route                  0
Rx Register no decap if                0
Null Register Timeout                 0
RP Filtered Source                    0
Rx Unknown Reg Stop                   0
Rx Join/Prune no state                 0
Rx Join/Prune on upstream if           0
Rx Join/Prune for invalid group        0
Rx Join/Prune messages dropped         0
Rx sparse join for dense group         0
Rx Graft/Graft Ack no state            0
Rx Graft on upstream if                0
Rx CRP not BSR                        0
Rx BSR when BSR                       0
Rx BSR not RPF if                     0
Rx unknown hello opt                   0
Rx data no state                       0

```


Rx RP no state	0
Rx aggregate	0
Rx malformed packet	0
Rx illegal TTL	0
Rx illegal destination address	0
No RP	0
No register encap if	0
No route upstream	28
Nexthop Unusable	0
RP mismatch	0
RP mode mismatch	0
RPF neighbor unknown	0
Rx Joins/Prunes filtered	0
Tx Joins/Prunes filtered	0
Embedded-RP invalid addr	0
Embedded-RP limit exceed	0
Embedded-RP added	0
Embedded-RP removed	0
Rx Register msgs filtering drop	0
Tx Register msgs filtering drop	0
Rx Bidir Join/Prune on non-Bidir if	0
Rx Bidir Join/Prune on non-DF if	0
V4 (S,G) Maximum	10
V4 (S,G) Accepted	9
V4 (S,G) Threshold	80
V4 (S,G) Log Interval	80
V6 (S,G) Maximum	8
V6 (S,G) Accepted	8
V6 (S,G) Threshold	50
V6 (S,G) Log Interval	100
V4 (grp-prefix, RP) Maximum	100
V4 (grp-prefix, RP) Accepted	5
V4 (grp-prefix, RP) Threshold	80
V4 (grp-prefix, RP) Log Interval	10
V6 (grp-prefix, RP) Maximum	20
V6 (grp-prefix, RP) Accepted	0
V6 (grp-prefix, RP) Threshold	90
V6 (grp-prefix, RP) Log Interval	20
V4 Register Maximum	100
V4 Register Accepted	10
V4 Register Threshold	80
V4 Register Log Interval	10
V6 Register Maximum	20
V6 Register Accepted	0
V6 Register Threshold	90
V6 Register Log Interval	20
(*,G) Join drop due to SSM range check	0

Sample Output

show pim statistics interface <interface-name>

```

user@host> show pim statistics interface ge-0/3/0.0
Instance: PIM.master Family: INET

PIM Interface statistics for ge-0/3/0.0

PIM Message type      Received      Sent  Rx errors
V2 Hello               0             3       0
V2 Register            0             0       0
V2 Register Stop       0             0       0

```

V2 Join Prune	0	0	0
V2 Bootstrap	0	0	0
V2 Assert	0	0	0
V2 Graft	0	0	0
V2 Graft Ack	0	0	0
V2 Candidate RP	0	0	0
V1 Query	0	0	0
V1 Register	0	0	0
V1 Register Stop	0	0	0
V1 Join Prune	0	0	0
V1 RP Reachability	0	0	0
V1 Assert	0	0	0
V1 Graft	0	0	0
V1 Graft Ack	0	0	0
AutoRP Announce	0	0	0
AutoRP Mapping	0	0	0
AutoRP Unknown type	0		
Anycast Register	0	0	0
Anycast Register Stop	0	0	0

Instance: PIM.master Family: INET6

PIM Interface statistics for ge-0/3/0.0

PIM Message type	Received	Sent	Rx errors
V2 Hello	0	3	0
V2 Register	0	0	0
V2 Register Stop	0	0	0
V2 Join Prune	0	0	0
V2 Bootstrap	0	0	0
V2 Assert	0	0	0
V2 Graft	0	0	0
V2 Graft Ack	0	0	0
V2 Candidate RP	0	0	0
Anycast Register	0	0	0
Anycast Register Stop	0	0	0

show sap listen

Syntax	show sap listen <brief detail> <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display the addresses that the router is listening to in order to receive multicast Session Announcement Protocol (SAP) session announcements.
Options	<p>none—Display standard information about the addresses that the router is listening to in order to receive multicast SAP session announcements.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show sap listen on page 585 show sap listen brief on page 585 show sap listen detail on page 586
Output Fields	Table 79 on page 585 describes the output fields for the show sap listen command. Output fields are listed in the approximate order in which they appear.

Table 79: show sap listen Output Fields

Field Name	Field Description
Group address	Address of the group that the local router is listening to for SAP messages.
Port	UDP port number used for SAP.

Sample Output

show sap listen

```
user@host> show sap listen
Group address  Port
224.2.127.254  9875
239.255.255.255 9875
```

show sap listen brief

The output for the **show sap listen brief** command is identical to that for the **show sap listen** command. For sample output, see [show sap listen on page 585](#).

`show sap listen detail`

The output for the **show sap listen detail** command is identical to that for the **show sap listen** command. For sample output, see [show sap listen on page 585](#).

test msdp

Syntax	test msdp (dependent-peers <i>prefix</i> rpf-peer <i>originator</i>) <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Find Multicast Source Discovery Protocol (MSDP) peers.
Options	<p>dependent-peers <i>prefix</i>—Find downstream dependent MSDP peers.</p> <p>rpf-peer <i>originator</i>—Find the MSDP reverse-path-forwarding (RPF) peer for the originator.</p> <p>instance <i>instance-name</i>—(Optional) Find MDSP peers for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	test msdp dependent-peers on page 587
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

test msdp dependent-peers

```
user@host> test msdp dependent-peers 10.0.0.1/24
```


CHAPTER 11

IPv6 Operational Mode Commands

- `clear ipv6 neighbors`
- `clear ipv6 router-advertisement`
- `show ipv6 neighbors`
- `show ipv6 router-advertisement`

clear ipv6 neighbors

Syntax	<code>clear ipv6 neighbors</code> <code><all host <i>hostname</i>></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.3 for EX Series switches. Command introduced in Junos OS Release 12.2 for the QFX Series.
Description	Clear IPv6 neighbor cache information.
Options	none —Clear all IPv6 neighbor cache information. all —(Optional) Clear all IPv6 neighbor cache information. host <i>hostname</i> —(Optional) Clear the information for the specified IPv6 neighbors.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show ipv6 neighbors on page 592
List of Sample Output	clear ipv6 neighbors on page 590
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ipv6 neighbors

```
user@host> clear ipv6 neighbors
```


clear ipv6 router-advertisement

Syntax	clear ipv6 router-advertisement <interface <i>interface</i> > <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Clear IPv6 router advertisement counters.
Options	<p>none—Clear IPv6 router advertisement counters for all interfaces.</p> <p>interface <i>interface</i>—(Optional) Clear IPv6 router advertisement counters for the specified interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show ipv6 router-advertisement on page 594
List of Sample Output	clear ipv6 router-advertisement on page 591
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ipv6 router-advertisement

```
user@host> clear ipv6 router-advertisement
```

show ipv6 neighbors

Syntax	show ipv6 neighbors
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.3 for EX Series switches. Command introduced in Junos OS Release 12.2 for the QFX Series.
Description	Display information about the IPv6 neighbor cache.
Options	This command has no options.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> clear ipv6 neighbors on page 590
List of Sample Output	show ipv6 neighbors on page 592
Output Fields	Table 80 on page 592 describes the output fields for the show ipv6 neighbors command. Output fields are listed in the approximate order in which they appear.

Table 80: show ipv6 neighbors Output Fields

Field Name	Field Description
IPv6 Address	Name of the IPv6 interface.
Linklayer Address	Link-layer address.
State	State of the link: up , down , incomplete , reachable , stale , or unreachable .
Exp	Number of seconds until the entry expires.
Rtr	Whether the neighbor is a routing device: yes or no .
Secure	Whether this entry was created using the Secure Neighbor Discovery (SEND) protocol: yes or no .
Interface	Name of the interface.

Sample Output

show ipv6 neighbors

```

user@host> show ipv6 neighbors
IPv6 Address          Linklayer Address  State      Exp Rtr Secure
Interface
2001:db8:0:1:2a0:a514:0:24c  00:05:85:8f:c8:bd  stale      546 yes no
fe-1/2/0.1

```

fe80::2a0:a514:0:24c fe-1/2/0.1	00:05:85:8f:c8:bd	stale	258	yes	no
fe80::2a0:a514:0:64c fe-1/2/1.5	00:05:85:8f:c8:bd	stale	111	yes	no
fe80::2a0:a514:0:a4c fe-1/2/2.9	00:05:85:8f:c8:bd	stale	327	yes	no

show ipv6 router-advertisement

Syntax	<pre>show ipv6 router-advertisement <conflicts> <interface <i>interface</i>> <logical-system (all <i>logical-system-name</i>)> <prefix <i>prefix/prefix length</i>></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 12.2 for the QFX Series.</p>
Description	Display information about IPv6 router advertisements, including statistics about messages sent and received on interfaces, and information received from advertisements from other routers.
Options	<p>none—Display all IPv6 router advertisement information for all interfaces.</p> <p>conflicts—(Optional) Display only the IPv6 router advertisement information that is conflicting.</p> <p>interface <i>interface</i>—(Optional) Display IPv6 router advertisement information for the specified interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>prefix <i>prefix/prefix length</i>—(Optional) Display IPv6 router advertisement information for the specified prefix.</p>
Additional Information	The display identifies conflicting information by enclosing the value the router is advertising in brackets.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> clear ipv6 router-advertisement on page 591
List of Sample Output	<p>show ipv6 router-advertisement on page 595</p> <p>show ipv6 router-advertisement conflicts on page 596</p> <p>show ipv6 router-advertisement prefix on page 596</p>
Output Fields	Table 81 on page 594 describes the output fields for the show ipv6 router-advertisement command. Output fields are listed in the approximate order in which they appear.

Table 81: show ipv6 router-advertisement Output Fields

Field Name	Field Description
Interface	Name of the interface.
Advertisements sent	Number of router advertisements sent and the elapsed time since they were sent.

Table 81: show ipv6 router-advertisement Output Fields (*continued*)

Field Name	Field Description
Solicits received	Number of solicitation messages received.
Advertisements received	Number of router advertisements received.
Advertisements from	Names of interfaces from which router advertisements have been received and the elapsed time since the last one was received.
Managed	Managed address configuration flag: 0 (stateless) or 1 (stateful).
Other configuration	Other stateful configuration flag: 0 (stateless) or 1 (stateful).
Reachable time	Time that a node identifies a neighbor as reachable after receiving a reachability confirmation, in milliseconds.
Default lifetime	Default lifetime, in seconds: from 0 seconds to 18.2 hours. A setting of 0 indicates that the router is not a default router.
Retransmit timer	Time between retransmitted Neighbor Solicitation messages, in milliseconds.
Current hop limit	Configured current hop limit.
Prefix	Name and length of the prefix.
Valid lifetime	How long the prefix remains valid for onlink determination.
Preferred lifetime	How long the prefix generated by stateless autoconfiguration remains preferred.
On link	Onlink flag: 0 (not onlink) or 1 (onlink).
Autonomous	Autonomous address configuration flag: 0 (not autonomous) or 1 (autonomous).

Sample Output

show ipv6 router-advertisement

```

user@host> show ipv6 router-advertisement
Interface: fe-0/1/1.0
  Advertisements sent: 0
  Solicits received: 0
  Advertisements received: 0
Interface: fxp0.0
  Advertisements sent: 0
  Solicits received: 0
  Advertisements received: 1
  Advertisement from fe80::2d0:b7ff:fe1e:7b0e, heard 00:00:13 ago
  Managed: 0
  Other configuration: 0 [1]
  Reachable time: 0 ms
  Default lifetime: 1800 sec

```

```
Retransmit timer: 0 ms
Current hop limit: 64
```

show ipv6 router-advertisement conflicts

```
user@host> show ipv6 router-advertisement conflicts
Interface: fxp0.0
  Advertisement from fe80::2d0:b7ff:fe1e:7b0e, heard 00:01:08 ago
  Other configuration: 0 [1]
```

show ipv6 router-advertisement prefix

```
user@host> show ipv6 router-advertisement prefix 8040::/16
Interface: fe-0/1/3.0
  Advertisements sent: 3, last sent 00:04:11 ago
  Solicits received: 0
  Advertisements received: 3
  Advertisement from fe80::290:69ff:fe9a:5403, heard 00:00:05 ago
  Managed: 0
  Other configuration: 0
  Reachable time: 0 ms
  Default lifetime: 180 sec [1800 sec]
  Retransmit timer: 0 ms
  Current hop limit: 64
  Prefix: 8040:1::/64
    Valid lifetime: 2592000 sec
    Preferred lifetime: 604800 sec
    On link: 1
    Autonomous: 1
```

CHAPTER 12

IS-IS Operational Mode Commands

- clear isis adjacency
- clear isis database
- clear isis overload
- clear isis statistics
- show isis adjacency
- show isis authentication
- show isis backup coverage
- show isis backup label-switched-path
- show isis backup spf results
- show isis context-identifier
- show isis database
- show isis hostname
- show isis interface
- show isis overview
- show isis route
- show isis spf
- show isis statistics

clear isis adjacency

Syntax	clear isis adjacency <instance <i>instance-name</i> > <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)> <neighbor>
Syntax (EX Series Switches and QFX Series)	clear isis adjacency <instance <i>instance-name</i> > <interface <i>interface-name</i> > <neighbor>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Remove entries from the IS-IS adjacency database.
Options	none —Remove all entries from the adjacency database. instance <i>instance-name</i> —(Optional) Clear all adjacencies for the specified routing instance only. interface <i>interface-name</i> —(Optional) Clear all adjacencies for the specified interface only. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. neighbor —(Optional) Clear adjacencies for the specified neighbor only.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show isis adjacency on page 606
List of Sample Output	clear isis adjacency on page 598
Output Fields	See show isis adjacency for an explanation of output fields.

Sample Output

clear isis adjacency

The following sample output displays IS-IS adjacency database information before and after the **clear isis adjacency** command is entered:

```
user@host> show isis adjacency
IS-IS adjacency database:
Interface      System          L State      Hold (secs) SNPA
so-1/0/0.0     karaku1         3 Up          26
so-1/1/3.0     1921.6800.5080 3 Up          23
```



```
so-5/0/0.0    1921.6800.5080 3 Up                                19
```

```
user@host> clear isis adjacency karakul
```

```
user@host> show isis adjacency
```

```
IS-IS adjacency database:
```

Interface	System	L State	Hold (secs)	SNPA
so-1/0/0.0	karakul	3 Initializing	26	
so-1/1/3.0	1921.6800.5080	3 Up	24	
so-5/0/0.0	1921.6800.5080	3 Up	21	

clear isis database

Syntax	clear isis database <entries> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and QFX Series)	clear isis database <entries> <instance <i>instance-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Remove the entries from the IS-IS link-state database, which contains prefixes and topology information.
Options	none —Remove all entries from the IS-IS link-state database for all routing instances. entries —(Optional) Name of the database entry. instance <i>instance-name</i> —(Optional) Clear all entries for the specified routing instance. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show isis database on page 622
List of Sample Output	clear isis database on page 600
Output Fields	See show isis database for an explanation of output fields.

Sample Output

clear isis database

The following sample output displays IS-IS link-state database information before and after the **clear isis database** command is entered:

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                Sequence Checksum Lifetime (secs)
crater.00-00          0x12   0x84dd             1139
  1 LSPs
IS-IS level 2 link-state database:
LSP ID                Sequence Checksum Lifetime (secs)
crater.00-00          0x19   0xe92c             1134
badlands.00-00        0x16   0x1454             985
carlsbad.00-00        0x33   0x220b             1015
```

ranier.00-00	0x2e	0xfc31	1007
1921.6800.5066.00-00	0x11	0x7313	566
1921.6800.5067.00-00	0x14	0xd9d4	939
6 LSPs			

user@host> clear isis database

user@host> show isis database

IS-IS level 1 link-state database:

LSP ID	Sequence	Checksum	Lifetime (secs)
--------	----------	----------	-----------------

IS-IS level 2 link-state database:

LSP ID	Sequence	Checksum	Lifetime (secs)
--------	----------	----------	-----------------

clear isis overload

Syntax	clear isis overload <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and QFX Series)	clear isis overload <instance <i>instance-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Reset the IS-IS dynamic overload bit. This command can appear to not work, continuing to display overload after execution. The bit is reset only if the root cause is corrected by configuration remotely or locally.</p> <p>When other routers detect that the overload bit is set, they do not use this routing device for transit traffic, but they do use it for packets destined to the overloaded routing device's directly connected networks and IP prefixes.</p>
Options	<p>none—Reset the IS-IS dynamic overload bit.</p> <p>instance <i>instance-name</i>—(Optional) Reset the IS-IS dynamic overload bit for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show isis database on page 622
List of Sample Output	clear isis overload on page 602
Output Fields	See show isis database for an explanation of output fields.

Sample Output

clear isis overload

The following sample output displays IS-IS database information before and after the **clear isis overload** command is entered:

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                Sequence Checksum Lifetime Attributes
pro3-c.00-00          0x4    0x10db    1185 L1 L2 Overload

  1 LSPs
```

IS-IS level 2 link-state database:

LSP ID	Sequence	Checksum	Lifetime	Attributes
pro3-c.00-00	0x5	0x429f	1185	L1 L2 Overload

pro2-a.00-00	0x91e	0x2589	874	L1 L2
pro2-a.02-00	0x1	0xcbc	874	L1 L2

3 LSPs

user@host> clear isis overload

user@host> show isis database

IS-IS level 1 link-state database:

LSP ID	Sequence	Checksum	Lifetime	Attributes
pro3-c.00-00	0xa	0x429e	1183	L1 L2

1 LSPs

IS-IS level 2 link-state database:

LSP ID	Sequence	Checksum	Lifetime	Attributes
pro3-c.00-00	0xc	0x9c39	1183	L1 L2
pro2-a.00-00	0x91e	0x2589	783	L1 L2
pro2-a.02-00	0x1	0xcbc	783	L1 L2

3 LSPs

clear isis statistics

Syntax	clear isis statistics <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and QFX Series)	clear isis statistics <instance <i>instance-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Set statistics about IS-IS traffic to zero.
Options	none —Set IS-IS traffic statistics to zero for all routing instances. instance <i>instance-name</i> —(Optional) Set IS-IS traffic statistics to zero for the specified routing instance only. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show isis statistics on page 652
List of Sample Output	clear isis statistics on page 604
Output Fields	See show isis statistics for an explanation of output fields.

Sample Output

clear isis statistics

The following sample output displays IS-IS statistics before and after the **clear isis statistics** command is entered:

```
user@host> show isis statistics
IS-IS statistics for merino:
```

PDU type	Received	Processed	Drops	Sent	Rexmit
LSP	12793	12793	0	8666	719
IIH	116751	116751	0	118834	0
CSNP	203956	203956	0	204080	0
PSNP	7356	7350	6	8635	0
Unknown	0	0	0	0	0
Totals	340856	340850	6	340215	719

Total packets received: 340856 Sent: 340934

SNP queue length: 0 Drops: 0

LSP queue length: 0 Drops: 0

SPF runs: 1064
Fragments rebuilt: 1087
LSP regenerations: 436
Purges initiated: 0

user@host> clear isis statistics

user@host> show isis statistics
IS-IS statistics for merino:

PDU type	Received	Processed	Drops	Sent	Rexmit
LSP	0	0	0	0	0
IIH	3	3	0	3	0
CSNP	2	2	0	4	0
PSNP	0	0	0	0	0
Unknown	0	0	0	0	0
Totals	5	5	0	7	0

Total packets received: 5 Sent: 7

SNP queue length: 0 Drops: 0
LSP queue length: 0 Drops: 0

SPF runs: 0
Fragments rebuilt: 0
LSP regenerations: 0
Purges initiated: 0

show isis adjacency

Syntax	show isis adjacency <system-id> <brief detail extensive> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>	
Syntax (EX Series Switches and QFX Series)	show isis adjacency <system-id> <brief detail extensive> <instance <i>instance-name</i> >	
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.	
Description	Display information about IS-IS neighbors.	
Options	<p>none—Display standard information about IS-IS neighbors for all routing instances.</p> <p>system id—(Optional) Display information about IS-IS neighbors for the specified intermediate system.</p> <p>brief detail extensive—(Optional) Display standard information about IS-IS neighbors with the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display information about IS-IS neighbors for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Display information about IS-IS neighbors for all logical systems or for a particular logical system.</p>	
Required Privilege Level	view	
Related Documentation	<ul style="list-style-type: none"> • clear isis adjacency on page 598 	
List of Sample Output	show isis adjacency on page 608 show isis adjacency brief on page 608 show isis adjacency detail on page 609 show isis adjacency extensive on page 609	
Output Fields	Table 82 on page 606 describes the output fields for the show isis adjacency command. Output fields are listed in the approximate order in which they appear.	

Table 82: show isis adjacency Output Fields

Field Name	Field Description	Level of Output
Interface	Interface through which the neighbor is reachable.	All levels

Table 82: show isis adjacency Output Fields (*continued*)

Field Name	Field Description	Level of Output
System	System identifier (sysid), displayed as a name, if possible.	brief
L or Level	Level: <ul style="list-style-type: none"> • 1—Level 1 only • 2—Level 2 only • 3—Level 1 and Level 2 An exclamation point (!) preceding the level number indicates that the adjacency is missing an IP address.	All levels
State	State of the adjacency: Up , Down , New , One-way , Initializing , or Rejected .	All levels
Hold (secs)	Remaining hold time of the adjacency.	brief
SNPA	Subnetwork point of attachment (MAC address of the next hop).	brief
Expires in	How long until the adjacency expires, in seconds.	detail
Priority	Priority to become the designated intermediate system.	detail extensive
Up/Down transitions	Count of adjacency status changes from Up to Down or from Down to Up .	detail
Last transition	Time of the last Up/Down transition.	detail
Circuit type	Bit mask of levels on this interface: 1=Level 1 router; 2=Level 2 router; 3=both Level 1 and Level 2 router.	detail
Speaks	Protocols supported by this neighbor.	detail extensive
MAC address	MAC address of the interface.	detail extensive
Topologies	Supported topologies.	detail extensive
Restart capable	Whether a neighbor is capable of graceful restart: Yes or No .	detail extensive
Adjacency advertisement: Advertise	This routing device has signaled to advertise this interface to its neighbors in their link-state PDUs.	detail extensive
Adjacency advertisement: Suppress	This neighbor has signaled not to advertise the interface in the routing device's outbound link-state PDUs.	detail extensive
IP addresses	IP address of this neighbor.	detail extensive

Table 82: show isis adjacency Output Fields (*continued*)

Field Name	Field Description	Level of Output
Transition log	<p>List of recent transitions, including:</p> <ul style="list-style-type: none"> • When—Time at which an IS-IS adjacency transition occurred. • State—Current state of the IS-IS adjacency (up, down, or rejected). <ul style="list-style-type: none"> • Up—Adjacency is up and operational. • Down—Adjacency is down and not available. • Rejected—Adjacency has been rejected. • Event—Type of transition that occurred. <ul style="list-style-type: none"> • Seenself—Possible routing loop has been detected. • Interface down—IS-IS interface has gone down and is no longer available. • Error—Adjacency error. • Down reason—Reason that an IS-IS adjacency is down: <ul style="list-style-type: none"> • 3-Way Handshake Failed—Connection establishment failed. • Address Mismatch—Address mismatch caused link failure. • Aged Out—Link expired. • ISO Area Mismatch—IS-IS area mismatch caused link failure. • Bad Hello—Unacceptable hello message caused link failure. • BFD Session Down—Bidirectional failure detection caused link failure. • Interface Disabled—IS-IS interface is disabled. • Interface Down—IS-IS interface is unavailable. • Interface Level Disabled—IS-IS level is disabled. • Level Changed—IS-IS level has changed on the adjacency. • Level Mismatch—Levels on adjacency are not compatible. • MPLS LSP Down—Label-switched path (LSP) is unavailable. • MT Topology Changed—IS-IS topology has changed. • MT Topology Mismatch—IS-IS topology is mismatched. • Remote System ID Changed—Adjacency peer system ID changed. • Protocol Shutdown—IS-IS protocol is disabled. • CLI Command—Adjacency brought down by user. • Unknown—Unknown. 	extensive

Sample Output

show isis adjacency

```

user@host> show isis adjacency
Interface          System      L State      Hold (secs) SNPA
at-2/3/0.0         ranier      3 Up         23

```

show isis adjacency brief

The output for the **show isis adjacency brief** command is identical to that for the **show isis adjacency** command. For sample output, see [show isis adjacency on page 608](#).

show isis adjacency detail

```
user@host> show isis adjacency detail
ranier
Interface: at-2/3/0.0, Level: 3, State: Up, Expires in 21 secs
Priority: 0, Up/Down transitions: 1, Last transition: 00:01:09 ago
Circuit type: 3, Speaks: IP, IPv6
Topologies: Unicast
Restart capable: Yes
IP addresses: 11.1.1.2
```

show isis adjacency extensive

```
user@host> show isis adjacency extensive
ranier
Interface: at-2/3/0.0, Level: 3, State: Up, Expires in 22 secs
Priority: 0, Up/Down transitions: 1, Last transition: 00:01:16 ago
Circuit type: 3, Speaks: IP, IPv6
Topologies: Unicast
Restart capable: Yes
IP addresses: 11.1.1.2
Transition log:
When          State      Event      Down reason
Wed Nov  8 21:24:25  Up        Seenself
```

show isis authentication

Syntax	show isis authentication <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and QFX Series)	show isis authentication <instance <i>instance-name</i> >
Release Information	Command introduced in Junos OS Release 7.5. Command introduced in Junos OS Release 9.0 for EX Series switches. Support for hitless authentication key rollover introduced in Junos OS Release 11.2. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display information about IS-IS authentication.
Options	<p>none—Display information about IS-IS authentication.</p> <p>instance <i>instance-name</i>—(Optional) Display IS-IS authentication for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show isis authentication on page 611 show isis authentication (With Hitless Authentication Key Rollover Configured) on page 611
Output Fields	Table 83 on page 610 describes the output fields for the show isis authentication command. Output fields are listed in the approximate order in which they appear.

Table 83: show isis authentication Output Fields

Field Name	Field Description
Interface	Interface name.
Level	IS-IS level.
IIH Auth	IS-IS Hello (IIH) packet authentication type. Displays the name of the active keychain if hitless authentication key rollover is configured.
CSN Auth	Complete sequence number authentication type.
PSN Auth	Partial sequence number authentication type.

Table 83: show isis authentication Output Fields *(continued)*

Field Name	Field Description
L1 LSP Authentication	Layer 1 link-state PDU authentication type.
L2 LSP Authentication	Layer 2 link-state PDU authentication type.

Sample Output

show isis authentication

```

user@host> show isis authentication
Interface          Level IIH Auth  CSN Auth  PSN Auth
at-2/3/0.0         1      Simple    Simple    Simple
                   2      MD5       MD5       MD5

L1 LSP Authentication: Simple
L2 LSP Authentication: MD5

```

show isis authentication (With Hitless Authentication Key Rollover Configured)

```

user@host> show isis authentication
Interface          Level IIH Auth  CSN Auth  PSN Auth
so-0/1/3.0         2      hakrhello MD5       MD5

L2 LSP Authentication: MD5

```

show isis backup coverage

Syntax	<pre>show isis backup coverage <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show isis backup coverage <instance <i>instance-name</i>></pre>
Release Information	<p>Command introduced in Junos OS Release 9.5.</p> <p>Command introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Display information about the level of backup coverage available.
Options	<p>none—Display information about the level of backup coverage available for all the nodes and prefixes in the network.</p> <p>instance <i>instance-name</i>—(Optional) Display information about the level of backup coverage for a specific IS-IS routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Link and Node Protection for IS-IS Routes</i> • show isis backup label-switched-path on page 614
List of Sample Output	show isis backup coverage on page 613
Output Fields	<p>Table 84 on page 612 lists the output fields for the show isis backup coverage command. Output fields are listed in the approximate order in which they appear.</p>

Table 84: show isis backup coverage Output Fields

Field Name	Field Description
Topology	Type of topology or address family: IPV4 Unicast or IPV6 Unicast .
Level	IS-IS level: <ul style="list-style-type: none"> • 1—Level 1 • 2—Level 2
Node	By topology, the percentage of all routes configured on the node that are protected through backup coverage.

Table 84: show isis backup coverage Output Fields (*continued*)

Field Name	Field Description
IPv4	Percentage of IPv4 unicast routes that are protected through backup coverage.
IPv6	Percentage of IPv6 unicast routes that are protected through backup coverage.
CLNS	Percentage of Connectionless Network Service (CLNS) routes that are protected through backup coverage.

Sample Output

show isis backup coverage

```

user@host> show isis backup coverage
Backup Coverage:
  Topology   Level  Node   IPv4   IPv6   CLNS
  IPV4 Unicast  2  28.57%  22.22%  0.00%  0.00%
  IPV6 Unicast  2   0.00%  0.00%  0.00%  0.00%

```

show isis backup label-switched-path

Syntax	show isis backup label-switched-path <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and QFX Series)	show isis backup label-switched-path
Release Information	Command introduced in Junos OS Release 9.5. Command introduced in Junos OS Release 9.5 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display information about MPLS label-switched-paths (LSPs) designated as backup routes for IS-IS routes.
Options	none —Display information about MPLS LSPs designated as backup routes for IS-IS routes. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Link and Node Protection for IS-IS Routes</i> • show isis backup coverage on page 612
List of Sample Output	show isis backup label-switched-path on page 615
Output Fields	Table 85 on page 614 lists the output fields for the show isis backup label-switched-path command. Output fields are listed in the approximate order in which they appear.

Table 85: show isis backup label-switched-path Output Fields

Field Name	Field Description
Backup MPLS LSPs	List of MPLS LSPs designated as backup paths for IS-IS routes.
Egress	IP address of the egress routing device for the LSP.
Status	State of the LSP: <ul style="list-style-type: none"> • Up—The routing device can detect RSVP hello messages from the neighbor. • Down—The routing device has received one of the following indications: <ul style="list-style-type: none"> • Communication failure from the neighbor. • Communication from IGP that the neighbor is unavailable. • Change in the sequence numbers in the RSVP hello messages sent by the neighbor. • Deleted—LSP is no longer available as a backup path.

Table 85: show isis backup label-switched-path Output Fields (*continued*)

Field Name	Field Description
Last change	Time elapsed since the neighbor state changed either from up to down or from down to up. The format is <i>hh:mm:ss</i> .
TE-metric	Configured traffic engineering metric.
Metric	Configured metric.

Sample Output

show isis backup label-switched-path

```
user@host> show isis backup label-switched-path
Backup MPLS LSPs:
f-to-g, Egress: 192.168.1.4, Status: up, Last change: 06:12:03
TE-metric: 9, Metric: 0
```

show isis backup spf results

Syntax	<code>show isis backup spf results</code> <code><instance <i>instance-name</i>></code> <code><level (1 2)></code> <code><logical-system (all <i>logical-system-name</i>)></code> <code><no-coverage></code> <code><topology (ipv4-unicast ipv6-multicast ipv6-unicast unicast)></code>
Syntax (EX Series Switches)	<code>show isis backup spf results</code> <code><instance <i>instance-name</i>></code> <code><level (1 2)></code> <code><no-coverage></code> <code><topology (ipv4-unicast unicast)></code>
Release Information	Command introduced in Junos OS Release 9.5.
Description	Display information about IS-IS shortest-path-first (SPF) calculations for backup paths.
Options	<p>none—Display information about IS-IS SPF calculations for all backup paths for all destination nodes.</p> <p>instance <i>instance-name</i>—(Optional) Display SPF calculations for backup paths for the specified routing instance.</p> <p>level (1 2)—(Optional) Display SPF calculations for the backup paths for the specified IS-IS level.</p> <p>logical-system <i>logical-system-name</i>—(Optional) Display SPF calculations for the backup paths for all logical systems or on a particular logical system.</p> <p>no-coverage—(Optional) Display SPF calculations only for destinations that do not have backup coverage.</p> <p>topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)—(Optional) Display SPF calculations for backup paths for the specified topology only.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring Link and Node Protection for IS-IS Routes</i>• show isis backup coverage on page 612
List of Sample Output	show isis backup spf results on page 617 show isis backup spf results no-coverage on page 618
Output Fields	Table 86 on page 617 lists the output fields for the show isis backup spf results command. Output fields are listed in the approximate order in which they appear.

Table 86: show isis backup spf results Output Fields

Field Name	Field Description
<i>node-name</i>	Name of the destination node.
Address	Address of the destination node.
Primary next-hop	Interface and name of the node of the primary next hop to reach the destination.
Root	Name of the next-hop neighbor.
Metric	Metric to the node.
Eligible	Indicates that the next-hop neighbor has been designated as a backup path to the destination node.
Backup next-hop	Name of the interface of the backup next hop.
SNPA	Subnetwork point of attachment (MAC address of the next hop).
LSP	Name of the MPLS label-switched path (LSP) designated as a backup path.
Not eligible	Indicates that the next-hop neighbor cannot function as a backup path to the destination.
Reason	Describes why the next-hop neighbor is designated as Not eligible as a backup path.

Sample Output

show isis backup spf results

```

user@host> show isis backup spf results
IS-IS level 1 SPF results:
pro-bng3-k.00
  Primary next-hop: fe-1/3/3.0, IPV4, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
  Primary next-hop: fe-1/3/3.0, IPV6, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
  Root: pro-bng3-k, Root Metric: 10, Metric: 0, Root Preference: 0x0
  Not eligible, IPV4, Reason: Primary next-hop link fate sharing
  Not eligible, IPV6, Reason: Primary next-hop link fate sharing
  Root: pro-bng3-i, Root Metric: 10, Metric: 20, Root Preference: 0x0
  track-item: pro-bng3-k.00-00
  track-item: pro-bng3-j.00-00
  Not eligible, IPV4, Reason: Path loops
  Not eligible, IPV6, Reason: Path loops
pro-bng3-i.00
  Primary next-hop: fe-0/1/2.0, IPV4, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21
  Primary next-hop: fe-0/1/2.0, IPV6, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21
  Root: pro-bng3-i, Root Metric: 10, Metric: 0, Root Preference: 0x0
  Not eligible, IPV4, Reason: Primary next-hop link fate sharing
  Not eligible, IPV6, Reason: Primary next-hop link fate sharing
  Root: pro-bng3-k, Root Metric: 10, Metric: 20, Root Preference: 0x0

```

```

        track-item: pro-bng3-j.00-00
        track-item: pro-bng3-i.00-00
        Not eligible, IPV4, Reason: Path loops
        Not eligible, IPV6, Reason: Path loops
    2 nodes

IS-IS level 2 SPF results:
pro-bng3-k.00
  Primary next-hop: fe-1/3/3.0, IPV4, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
  Primary next-hop: fe-1/3/3.0, IPV6, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
  Root: pro-bng3-k, Root Metric: 10, Metric: 0, Root Preference: 0x0
  Not eligible, IPV4, Reason: Primary next-hop link fate sharing
  Not eligible, IPV6, Reason: Primary next-hop link fate sharing
  Root: pro-bng3-i, Root Metric: 10, Metric: 20, Root Preference: 0x0
  track-item: pro-bng3-k.00-00
  track-item: pro-bng3-j.00-00
  Not eligible, IPV4, Reason: Path loops
  Not eligible, IPV6, Reason: Path loops
pro-bng3-i.00
  Primary next-hop: fe-0/1/2.0, IPV4, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21
  Primary next-hop: fe-0/1/2.0, IPV6, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21
  Root: pro-bng3-i, Root Metric: 10, Metric: 0, Root Preference: 0x0
  Not eligible, IPV4, Reason: Primary next-hop link fate sharing
  Not eligible, IPV6, Reason: Primary next-hop link fate sharing
  Root: pro-bng3-k, Root Metric: 10, Metric: 20, Root Preference: 0x0
  track-item: pro-bng3-j.00-00
  track-item: pro-bng3-i.00-00
  Not eligible, IPV4, Reason: Path loops
  Not eligible, IPV6, Reason: Path loops
2 nodes

```

show isis backup spf results no-coverage

```

user@host> show isis backup spf results no-coverage
IS-IS level 1 SPF results:
pro-bng3-k.00
  Primary next-hop: fe-1/3/3.0, IPV4, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
  Primary next-hop: fe-1/3/3.0, IPV6, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
  Root: pro-bng3-k, Root Metric: 10, Metric: 0, Root Preference: 0x0
  Root: pro-bng3-i, Root Metric: 10, Metric: 20, Root Preference: 0x0
  track-item: pro-bng3-k.00-00
  track-item: pro-bng3-j.00-00
pro-bng3-i.00
  Primary next-hop: fe-0/1/2.0, IPV4, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21
  Primary next-hop: fe-0/1/2.0, IPV6, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21
  Root: pro-bng3-i, Root Metric: 10, Metric: 0, Root Preference: 0x0
  Root: pro-bng3-k, Root Metric: 10, Metric: 20, Root Preference: 0x0
  track-item: pro-bng3-j.00-00
  track-item: pro-bng3-i.00-00
2 nodes

IS-IS level 2 SPF results:
pro-bng3-k.00
  Primary next-hop: fe-1/3/3.0, IPV4, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
  Primary next-hop: fe-1/3/3.0, IPV6, pro-bng3-k, SNPA: b0:c6:9a:2c:f0:de
  Root: pro-bng3-k, Root Metric: 10, Metric: 0, Root Preference: 0x0
  Root: pro-bng3-i, Root Metric: 10, Metric: 20, Root Preference: 0x0
  track-item: pro-bng3-k.00-00
  track-item: pro-bng3-j.00-00
pro-bng3-i.00
  Primary next-hop: fe-0/1/2.0, IPV4, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21

```

```
Primary next-hop: fe-0/1/2.0, IPV6, pro-bng3-i, SNPA: b0:c6:9a:2a:f4:21
Root: pro-bng3-i, Root Metric: 10, Metric: 0, Root Preference: 0x0
Root: pro-bng3-k, Root Metric: 10, Metric: 20, Root Preference: 0x0
  track-item: pro-bng3-j.00-00
  track-item: pro-bng3-i.00-00
2 nodes
```

show isis context-identifier

Syntax	show isis context-identifier <brief detail extensive> <identifier name> <instance instance-name> <logical-system (all logical-system-name)>
Release Information	Command introduced in Junos OS Release 10.4.
Description	Display IS-IS context identifier information.
Options	brief detail extensive —(Optional) Display the specified level of output. identifier name —(Optional) Display information about the specified context identifier. instance instance-name —(Optional) Display entries for the specified routing instance. logical-system (all logical-system-name) —(Optional) Display the context identifier information for all logical systems or for a particular logical system.
Required Privilege Level	View
Output Fields	Table 87 on page 620 lists the output fields for the show isis context-identifier command. Output fields are listed in the approximate order in which they appear.

Table 87: show isis context-identifier Output Fields

Field Name	Field Description	Level of Output
Context	IPv4 address that defines a protection pair. The context is manually configured on both primary and protector PEs.	detail
Owner	Protocol that requires the context.	detail
Role	Role of the PE, which is either primary or protector.	detail
Primary	Name of the primary PE.	detail
Metric	Advertised interior gateway protocol (IGP) metric.	detail

Sample Output

```
user@host> show isis context-identifier detail
```

```
IS-IS context database:
```

```
Context      Owner      Role      Primary      Metric
2.2.4.3      MPLS      Primary   pro3-e        1
  Advertiser pro3-e, Router ID 10.255.245.198, Metric 1, Level 1
  Advertiser pro3-e, Router ID 10.255.245.198, Metric 1, Level 2
  Advertiser pro3-c, Router ID 10.255.245.196, Metric 11, Level 2
```


show isis database

Syntax	<code>show isis database</code> <code><system-id></code> <code><brief detail extensive></code> <code><instance <i>instance-name</i>></code> <code><level (1 2)></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches and QFX Series)	<code>show isis database</code> <code><system-id></code> <code><brief detail extensive></code> <code><level (1 2)></code> <code><instance <i>instance-name</i>></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display the entries in the IS-IS link-state database, which contains data about PDU packets.
Options	<p>none—Display standard information about IS-IS link-state database entries for all routing instances.</p> <p><i>system id</i>—(Optional) Display IS-IS link-state database entries for the specified intermediate system.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display IS-IS link-state database entries for the specified routing instance.</p> <p>level (1 2)—(Optional) Display IS-IS link-state database entries for the specified IS-IS level.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Display standard information about IS-IS link-state database entries for all logical systems or for a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• clear isis database on page 600
List of Sample Output	show isis database on page 624 show isis database brief on page 625 show isis database detail on page 625 show isis database extensive on page 625

Output Fields Table 88 on page 623 describes the output fields for the **show isis database** command. Output fields are listed in the approximate order in which they appear. Fields that contain internal IS-IS information useful only in troubleshooting obscure problems are not described in the table. For more details about these fields, contact your customer support representative.

Table 88: show isis database Output Fields

Field Name	Field Description	Level of Output
Interface name	Name of the interface on which the link-state PDU has been received; always IS-IS for this command.	All levels
level	Level of intermediate system: <ul style="list-style-type: none"> • 1—Intermediate system routes within an area; when the destination is outside an area, it routes toward a Level 2 system. • 2—Intermediate system routes between areas and toward other ASs. 	All levels
LSP ID	Link-state PDU identifier.	All levels
Sequence	Sequence number of the link-state PDU.	All levels
Checksum	Checksum value of the link-state PDU.	All levels
Lifetime (secs)	Remaining lifetime of the link-state PDU, in seconds.	All levels
Attributes	Attributes of the specified database: L1 , L2 , Overload , or Attached (L1 only).	none brief
# LSPs	Total number of link-state PDUs in the specified link-state database.	none brief
IP prefix	Prefix advertised by this link-state PDU.	detail extensive
IS neighbor	IS-IS neighbor of the advertising system.	detail extensive
ES neighbor	(J Series routers only) An ES-IS neighbor of the advertising system.	detail extensive
IP prefix	IPv4 prefix advertised by this link-state PDU.	detail extensive
V6 prefix	IPv6 prefix advertised by this link-state PDU.	detail extensive
Metric	Metric of the prefix or neighbor.	detail extensive
Header	<ul style="list-style-type: none"> • LSP ID—Link state PDU identifier of the header. • Length—Header length. • Allocated Length—Amount of length available for the header. • Router ID—Address of the local routing device. • Remaining Lifetime—Remaining lifetime of the link-state PDU, in seconds. 	extensive

Table 88: show isis database Output Fields (*continued*)

Field Name	Field Description	Level of Output
Packet	<ul style="list-style-type: none"> • LSP ID—The identifier for the link-state PDU. • Length—Packet length. • Lifetime—Remaining lifetime, in seconds. • Checksum—The checksum of the link-state PDU. • Sequence—The sequence number of the link-state PDU. Every time the link-state PDU is updated, this number increments. • Attributes—Packet attributes. • NLPID—Network layer protocol identifier. • Fixed length—Specifies the set length for the packet. 	extensive
TLVs	<ul style="list-style-type: none"> • Area Address—Area addresses that the routing device can reach. • Speaks—Supported routing protocols. • IP router id—ID of the routing device (usually the IP address). • IP address—IPv4 address. • Hostname—Assigned name of the routing device. • IP prefix—IP prefix of the routing device. • Metric—IS-IS metric that measures the cost of the adjacency between the originating routing device and the advertised routing device. • IP extended prefix—Extended IP prefix of the routing device. • IS neighbor—Directly attached neighbor's name and metric. • IS extended neighbor—Directly attached neighbor's name, metric, IP address, local interface index, and remote interface index. <p>The interface indexes enable Junos OS to support unnumbered extensions for IS-IS, as described in RFC 4205.</p>	extensive

Sample Output

show isis database

```

user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                Sequence Checksum Lifetime Attributes
kobuk.00-00           0x3    0x3167    1057 L1 L2
camaro.00-00          0x5    0x770e    1091 L1 L2
ranier.00-00          0x4    0xaa95    1091 L1 L2
glacier.00-00         0x4    0x206f    1089 L1 L2
glacier.02-00         0x1    0xd141    1089 L1 L2
badlands.00-00        0x3    0x87a2    1093 L1 L2
  6 LSPs

IS-IS level 2 link-state database:
LSP ID                Sequence Checksum Lifetime Attributes
kobuk.00-00           0x6    0x8d6b    1096 L1 L2
camaro.00-00          0x9    0x877b    1101 L1 L2
ranier.00-00          0x8    0x855d    1103 L1 L2
glacier.00-00         0x7    0xf892    1098 L1 L2
glacier.02-00         0x1    0xd141    1089 L1 L2
badlands.00-00        0x6    0x562     1105 L1 L2
  6 LSPs

```

show isis database brief

The output for the **show isis database brief** command is identical to that for the **show isis database** command. For sample output, see [show isis database on page 624](#).

show isis database detail

```
user@host> show isis database logical-system CE3 sisira.00-00 detail
```

IS-IS level 1 link-state database:

```
sisira.00-00 Sequence: 0x11, Checksum: 0x10fc, Lifetime: 975 secs
  IS neighbor: hemantha-CE3.02           Metric:      10
  ES neighbor: 0015.0015.0015           Metric:      10 Down
  ES neighbor: 0025.0025.0025           Metric:      10 Down
  ES neighbor: 0030.0030.0030           Metric:      10 Down
  ES neighbor: 0040.0040.0040           Metric:      10 Down
  ES neighbor: sisira                     Metric:       0
  IP prefix: 1.0.0.0/24                  Metric:      10 External Down
  IP prefix: 3.0.0.0/24                  Metric:      10 External Down
  IP prefix: 4.0.0.0/24                  Metric:      10 External Down
  IP prefix: 5.0.0.0/24                  Metric:      10 Internal Up
  IP prefix: 15.15.15.15/32              Metric:      10 External Down
  IP prefix: 25.25.25.25/32              Metric:      10 External Down
  IP prefix: 30.30.30.30/32              Metric:      10 External Down
  IP prefix: 40.40.40.40/32              Metric:      10 External Down
  IP prefix: 60.60.60.60/32              Metric:       0 Internal Up
```

IS-IS level 2 link-state database:

```
sisira.00-00 Sequence: 0x13, Checksum: 0x69ac, Lifetime: 993 secs
  IS neighbor: hemantha-CE3.02           Metric:      10
  IP prefix: 1.0.0.0/24                  Metric:      10 External Down
  IP prefix: 3.0.0.0/24                  Metric:      10 External Down
  IP prefix: 4.0.0.0/24                  Metric:      10 External Down
  IP prefix: 5.0.0.0/24                  Metric:      10 Internal Up
  IP prefix: 15.15.15.15/32              Metric:      10 External Down
  IP prefix: 25.25.25.25/32              Metric:      10 External Down
  IP prefix: 30.30.30.30/32              Metric:      10 External Down
  IP prefix: 40.40.40.40/32              Metric:      10 External Down
  IP prefix: 50.50.50.50/32              Metric:      10 Internal Up
  IP prefix: 60.60.60.60/32              Metric:       0 Internal Up
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0015.0015.0015/152
                                          Metric:      10 External Down
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0025.0025.0025/152
                                          Metric:      10 External Down
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0030.0030.0030/152
                                          Metric:      10 External Down
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0040.0040.0040/152
                                          Metric:      10 External Down
  ISO prefix: 60.0006.80ff.f800.0000.0108.0001.0060.0060.0060/152
                                          Metric:       0 Internal Up
```

show isis database extensive

```
user@host> show isis database extensive
```

IS-IS level 1 link-state database:

```
Router-A.00-00 Sequence: 0x1, Checksum: 0xf75c, Lifetime: 1116 secs
```

IP prefix: 192.168.0.1/32 Metric: 0 Internal Up

Header: LSP ID: Router-A.00-00, Length: 85 bytes
Allocated length: 1492 bytes, Router ID: 192.168.0.1
Remaining lifetime: 1116 secs, Level: 1, Interface: 0
Estimated free bytes: 1353, Actual free bytes: 1407
Aging timer expires in: 1116 secs
Protocols: IP, IPv6

Packet: LSP ID: Router-A.00-00, Length: 85 bytes, Lifetime : 1200 secs
Checksum: 0xf75c, Sequence: 0x1, Attributes: 0x3 <L1 L2>
NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
Packet type: 18, Packet version: 1, Max area: 0

TLVs:

Area address: 49.0002 (3)
LSP Buffer Size: 1492
Speaks: IP
Speaks: IPV6
IP router id: 192.168.0.1
IP address: 192.168.0.1
Hostname: Router-A
IP prefix: 192.168.0.1/32, Internal, Metric: default 0, Up
IP extended prefix: 192.168.0.1/32 metric 0 up
No queued transmissions

IS-IS level 2 link-state database:

Router-A.00-00 Sequence: 0x5, Checksum: 0x3196, Lifetime: 1158 secs
IS neighbor: Router-B.02 Metric: 10
Two-way fragment: Router-B.02-00, Two-way first fragment: Router-B.02-00
IS neighbor: Router-E.02 Metric: 10
Two-way fragment: Router-E.02-00, Two-way first fragment: Router-E.02-00
IP prefix: 10.0.0.0/30 Metric: 10 Internal Up
IP prefix: 10.0.0.4/30 Metric: 10 Internal Up
IP prefix: 192.168.0.1/32 Metric: 0 Internal Up

Header: LSP ID: Router-A.00-00, Length: 208 bytes
Allocated length: 1492 bytes, Router ID: 192.168.0.1
Remaining lifetime: 1158 secs, Level: 2, Interface: 0
Estimated free bytes: 1233, Actual free bytes: 1284
Aging timer expires in: 1158 secs
Protocols: IP, IPv6

Packet: LSP ID: Router-A.00-00, Length: 208 bytes, Lifetime : 1198 secs
Checksum: 0x3196, Sequence: 0x5, Attributes: 0x3 <L1 L2>
NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
Packet type: 20, Packet version: 1, Max area: 0

TLVs:

Area address: 49.0002 (3)
LSP Buffer Size: 1492
Speaks: IP
Speaks: IPV6
IP router id: 192.168.0.1
IP address: 192.168.0.1
Hostname: Router-A
IP prefix: 192.168.0.1/32, Internal, Metric: default 0, Up
IP prefix: 10.0.0.4/30, Internal, Metric: default 10, Up
IP prefix: 10.0.0.0/30, Internal, Metric: default 10, Up
IP extended prefix: 192.168.0.1/32 metric 0 up

```

IP extended prefix: 10.0.0.4/30 metric 10 up
IP extended prefix: 10.0.0.0/30 metric 10 up
IS neighbor: Router-E.02, Internal, Metric: default 10
IS neighbor: Router-B.02, Internal, Metric: default 10
IS extended neighbor: Router-E.02, Metric: default 10
  IP address: 10.0.0.1
    Local interface index: 101, Remote interface index: 0
IS extended neighbor: Router-B.02, Metric: default 10
  IP address: 10.0.0.5
    Local interface index: 102, Remote interface index: 0
No queued transmissions

Router-B.00-00 Sequence: 0x5, Checksum: 0xf8f, Lifetime: 1183 secs
  IS neighbor: Router-B.02                      Metric: 10
    Two-way fragment: Router-B.02-00, Two-way first fragment: Router-B.02-00
  IS neighbor: Router-C.02                      Metric: 10
    Two-way fragment: Router-C.02-00, Two-way first fragment: Router-C.02-00
IP prefix: 10.0.0.4/30                          Metric: 10 Internal Up
IP prefix: 10.0.0.8/30                          Metric: 10 Internal Up
IP prefix: 192.168.0.2/32                      Metric: 0 Internal Up

Header: LSP ID: Router-B.00-00, Length: 208 bytes
  Allocated length: 284 bytes, Router ID: 192.168.0.2
  Remaining lifetime: 1183 secs, Level: 2, Interface: 102
  Estimated free bytes: 114, Actual free bytes: 76
  Aging timer expires in: 1183 secs
  Protocols: IP, IPv6

Packet: LSP ID: Router-B.00-00, Length: 208 bytes, Lifetime : 1196 secs
  Checksum: 0xf8f, Sequence: 0x5, Attributes: 0x3 <L1 L2>
  NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
  Packet type: 20, Packet version: 1, Max area: 0

TLVs:
  Area address: 49.0002 (3)
  LSP Buffer Size: 1492
  Speaks: IP
  Speaks: IPV6
  IP router id: 192.168.0.2
  IP address: 192.168.0.2
  Hostname: Router-B
  IP prefix: 192.168.0.2/32, Internal, Metric: default 0, Up
  IP prefix: 10.0.0.4/30, Internal, Metric: default 10, Up
  IP prefix: 10.0.0.8/30, Internal, Metric: default 10, Up
  IP extended prefix: 192.168.0.2/32 metric 0 up
  IP extended prefix: 10.0.0.4/30 metric 10 up
  IP extended prefix: 10.0.0.8/30 metric 10 up
  IS neighbor: Router-B.02, Internal, Metric: default 10
  IS neighbor: Router-C.02, Internal, Metric: default 10
  IS extended neighbor: Router-B.02, Metric: default 10
    IP address: 10.0.0.6
      Local interface index: 108, Remote interface index: 0
  IS extended neighbor: Router-C.02, Metric: default 10
    IP address: 10.0.0.9
      Local interface index: 109, Remote interface index: 0
No queued transmissions

Router-B.02-00 Sequence: 0x1, Checksum: 0x3c7c, Lifetime: 1156 secs
  IS neighbor: Router-A.00                      Metric: 0
    Two-way fragment: Router-A.00-00, Two-way first fragment: Router-A.00-00
  IS neighbor: Router-B.00                      Metric: 0

```

Two-way fragment: Router-B.00-00, Two-way first fragment: Router-B.00-00

Header: LSP ID: Router-B.02-00, Length: 76 bytes
Allocated length: 284 bytes, Router ID: 0.0.0.0
Remaining lifetime: 1156 secs, Level: 2, Interface: 102
Estimated free bytes: 208, Actual free bytes: 208
Aging timer expires in: 1156 secs

Packet: LSP ID: Router-B.02-00, Length: 76 bytes, Lifetime : 1196 secs
Checksum: 0x3c7c, Sequence: 0x1, Attributes: 0x3 <L1 L2>
NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
Packet type: 20, Packet version: 1, Max area: 0

TLVs:

IS neighbor: Router-B.00, Internal, Metric: default 0
IS neighbor: Router-A.00, Internal, Metric: default 0
IS extended neighbor: Router-B.00, Metric: default 0
IS extended neighbor: Router-A.00, Metric: default 0

No queued transmissions

Router-C.00-00 Sequence: 0x5, Checksum: 0x255b, Lifetime: 1182 secs
IS neighbor: Router-C.02 Metric: 10
Two-way fragment: Router-C.02-00, Two-way first fragment: Router-C.02-00
IS neighbor: Router-D.03 Metric: 10
Two-way fragment: Router-D.03-00, Two-way first fragment: Router-D.03-00
IP prefix: 10.0.0.8/30 Metric: 10 Internal Up
IP prefix: 10.0.0.12/30 Metric: 10 Internal Up
IP prefix: 192.168.0.3/32 Metric: 0 Internal Up

Header: LSP ID: Router-C.00-00, Length: 208 bytes
Allocated length: 284 bytes, Router ID: 192.168.0.3
Remaining lifetime: 1182 secs, Level: 2, Interface: 102
Estimated free bytes: 114, Actual free bytes: 76
Aging timer expires in: 1182 secs
Protocols: IP, IPv6

Packet: LSP ID: Router-C.00-00, Length: 208 bytes, Lifetime : 1196 secs
Checksum: 0x255b, Sequence: 0x5, Attributes: 0x3 <L1 L2>
NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
Packet type: 20, Packet version: 1, Max area: 0

TLVs:

Area address: 49.0002 (3)
LSP Buffer Size: 1492
Speaks: IP
Speaks: IPV6
IP router id: 192.168.0.3
IP address: 192.168.0.3
Hostname: Router-C
IP prefix: 192.168.0.3/32, Internal, Metric: default 0, Up
IP prefix: 10.0.0.8/30, Internal, Metric: default 10, Up
IP prefix: 10.0.0.12/30, Internal, Metric: default 10, Up
IP extended prefix: 192.168.0.3/32 metric 0 up
IP extended prefix: 10.0.0.8/30 metric 10 up
IP extended prefix: 10.0.0.12/30 metric 10 up
IS neighbor: Router-C.02, Internal, Metric: default 10
IS neighbor: Router-D.03, Internal, Metric: default 10
IS extended neighbor: Router-C.02, Metric: default 10
IP address: 10.0.0.10
Local interface index: 105, Remote interface index: 0
IS extended neighbor: Router-D.03, Metric: default 10

IP address: 10.0.0.13
 Local interface index: 106, Remote interface index: 0
 No queued transmissions

Router-C.02-00 Sequence: 0x1, Checksum: 0xaa09, Lifetime: 1181 secs
 IS neighbor: Router-B.00 Metric: 0
 Two-way fragment: Router-B.00-00, Two-way first fragment: Router-B.00-00
 IS neighbor: Router-C.00 Metric: 0
 Two-way fragment: Router-C.00-00, Two-way first fragment: Router-C.00-00

Header: LSP ID: Router-C.02-00, Length: 76 bytes
 Allocated length: 284 bytes, Router ID: 0.0.0.0
 Remaining lifetime: 1181 secs, Level: 2, Interface: 102
 Estimated free bytes: 208, Actual free bytes: 208
 Aging timer expires in: 1181 secs

Packet: LSP ID: Router-C.02-00, Length: 76 bytes, Lifetime : 1194 secs
 Checksum: 0xaa09, Sequence: 0x1, Attributes: 0x3 <L1 L2>
 NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
 Packet type: 20, Packet version: 1, Max area: 0

TLVs:
 IS neighbor: Router-C.00, Internal, Metric: default 0
 IS neighbor: Router-B.00, Internal, Metric: default 0
 IS extended neighbor: Router-C.00, Metric: default 0
 IS extended neighbor: Router-B.00, Metric: default 0
 No queued transmissions

Router-D.00-00 Sequence: 0x4, Checksum: 0x8ab7, Lifetime: 1180 secs
 IS neighbor: Router-D.02 Metric: 10
 Two-way fragment: Router-D.02-00, Two-way first fragment: Router-D.02-00
 IS neighbor: Router-D.03 Metric: 10
 Two-way fragment: Router-D.03-00, Two-way first fragment: Router-D.03-00
 IP prefix: 10.0.0.12/30 Metric: 10 Internal Up
 IP prefix: 10.0.0.20/30 Metric: 10 Internal Up
 IP prefix: 192.168.0.4/32 Metric: 0 Internal Up

Header: LSP ID: Router-D.00-00, Length: 208 bytes
 Allocated length: 284 bytes, Router ID: 192.168.0.4
 Remaining lifetime: 1180 secs, Level: 2, Interface: 102
 Estimated free bytes: 114, Actual free bytes: 76
 Aging timer expires in: 1180 secs
 Protocols: IP, IPv6

Packet: LSP ID: Router-D.00-00, Length: 208 bytes, Lifetime : 1192 secs
 Checksum: 0x8ab7, Sequence: 0x4, Attributes: 0x3 <L1 L2>
 NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
 Packet type: 20, Packet version: 1, Max area: 0

TLVs:
 Area address: 49.0002 (3)
 LSP Buffer Size: 1492
 Speaks: IP
 Speaks: IPV6
 IP router id: 192.168.0.4
 IP address: 192.168.0.4
 Hostname: Router-D
 IP prefix: 192.168.0.4/32, Internal, Metric: default 0, Up
 IP prefix: 10.0.0.12/30, Internal, Metric: default 10, Up
 IP prefix: 10.0.0.20/30, Internal, Metric: default 10, Up
 IP extended prefix: 192.168.0.4/32 metric 0 up

```
IP extended prefix: 10.0.0.12/30 metric 10 up
IP extended prefix: 10.0.0.20/30 metric 10 up
IS neighbor: Router-D.02, Internal, Metric: default 10
IS neighbor: Router-D.03, Internal, Metric: default 10
IS extended neighbor: Router-D.02, Metric: default 10
  IP address: 10.0.0.22
  Local interface index: 115, Remote interface index: 0
IS extended neighbor: Router-D.03, Metric: default 10
  IP address: 10.0.0.14
  Local interface index: 114, Remote interface index: 0
No queued transmissions

Router-D.02-00 Sequence: 0x1, Checksum: 0xebbc, Lifetime: 1128 secs
IS neighbor: Router-D.00 Metric: 0
  Two-way fragment: Router-D.00-00, Two-way first fragment: Router-D.00-00
IS neighbor: Router-F.00 Metric: 0
  Two-way fragment: Router-F.00-00, Two-way first fragment: Router-F.00-00

Header: LSP ID: Router-D.02-00, Length: 76 bytes
  Allocated length: 284 bytes, Router ID: 0.0.0.0
  Remaining lifetime: 1128 secs, Level: 2, Interface: 101
  Estimated free bytes: 208, Actual free bytes: 208
  Aging timer expires in: 1128 secs

Packet: LSP ID: Router-D.02-00, Length: 76 bytes, Lifetime : 1160 secs
  Checksum: 0xebbc, Sequence: 0x1, Attributes: 0x3 <L1 L2>
  NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
  Packet type: 20, Packet version: 1, Max area: 0

TLVs:
  IS neighbor: Router-D.00, Internal, Metric: default 0
  IS neighbor: Router-F.00, Internal, Metric: default 0
  IS extended neighbor: Router-D.00, Metric: default 0
  IS extended neighbor: Router-F.00, Metric: default 0
No queued transmissions

Router-D.03-00 Sequence: 0x1, Checksum: 0x129b, Lifetime: 1180 secs
IS neighbor: Router-C.00 Metric: 0
  Two-way fragment: Router-C.00-00, Two-way first fragment: Router-C.00-00
IS neighbor: Router-D.00 Metric: 0
  Two-way fragment: Router-D.00-00, Two-way first fragment: Router-D.00-00

Header: LSP ID: Router-D.03-00, Length: 76 bytes
  Allocated length: 284 bytes, Router ID: 0.0.0.0
  Remaining lifetime: 1180 secs, Level: 2, Interface: 101
  Estimated free bytes: 208, Actual free bytes: 208
  Aging timer expires in: 1180 secs

Packet: LSP ID: Router-D.03-00, Length: 76 bytes, Lifetime : 1192 secs
  Checksum: 0x129b, Sequence: 0x1, Attributes: 0x3 <L1 L2>
  NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
  Packet type: 20, Packet version: 1, Max area: 0

TLVs:
  IS neighbor: Router-D.00, Internal, Metric: default 0
  IS neighbor: Router-C.00, Internal, Metric: default 0
  IS extended neighbor: Router-D.00, Metric: default 0
  IS extended neighbor: Router-C.00, Metric: default 0
No queued transmissions

Router-E.00-00 Sequence: 0x4, Checksum: 0x9da9, Lifetime: 1155 secs
```



```

IS neighbor: Router-E.02                      Metric:      10
  Two-way fragment: Router-E.02-00, Two-way first fragment: Router-E.02-00
IS neighbor: Router-F.02                      Metric:      20
  Two-way fragment: Router-F.02-00, Two-way first fragment: Router-F.02-00
IP prefix: 10.0.0.0/30                        Metric:      10 Internal Up
IP prefix: 10.0.0.16/30                      Metric:      20 Internal Up
IP prefix: 192.168.0.5/32                    Metric:       0 Internal Up

```

```

Header: LSP ID: Router-E.00-00, Length: 208 bytes
  Allocated length: 284 bytes, Router ID: 192.168.0.5
  Remaining lifetime: 1155 secs, Level: 2, Interface: 101
  Estimated free bytes: 114, Actual free bytes: 76
  Aging timer expires in: 1155 secs
  Protocols: IP, IPv6

```

```

Packet: LSP ID: Router-E.00-00, Length: 208 bytes, Lifetime : 1185 secs
  Checksum: 0x9da9, Sequence: 0x4, Attributes: 0x3 <L1 L2>
  NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
  Packet type: 20, Packet version: 1, Max area: 0

```

TLVs:

```

Area address: 49.0002 (3)
LSP Buffer Size: 1492
Speaks: IP
Speaks: IPv6
IP router id: 192.168.0.5
IP address: 192.168.0.5
Hostname: Router-E
IP prefix: 192.168.0.5/32, Internal, Metric: default 0, Up
IP prefix: 10.0.0.16/30, Internal, Metric: default 20, Up
IP prefix: 10.0.0.0/30, Internal, Metric: default 10, Up
IP extended prefix: 192.168.0.5/32 metric 0 up
IP extended prefix: 10.0.0.16/30 metric 20 up
IP extended prefix: 10.0.0.0/30 metric 10 up
IS neighbor: Router-E.02, Internal, Metric: default 10
IS neighbor: Router-F.02, Internal, Metric: default 20
IS extended neighbor: Router-E.02, Metric: default 10
  IP address: 10.0.0.2
  Local interface index: 112, Remote interface index: 0
IS extended neighbor: Router-F.02, Metric: default 20
  IP address: 10.0.0.17
  Local interface index: 111, Remote interface index: 0
No queued transmissions

```

```

Router-E.02-00 Sequence: 0x1, Checksum: 0xb4fa, Lifetime: 1130 secs
IS neighbor: Router-A.00                      Metric:       0
  Two-way fragment: Router-A.00-00, Two-way first fragment: Router-A.00-00
IS neighbor: Router-E.00                      Metric:       0
  Two-way fragment: Router-E.00-00, Two-way first fragment: Router-E.00-00

```

```

Header: LSP ID: Router-E.02-00, Length: 76 bytes
  Allocated length: 284 bytes, Router ID: 0.0.0.0
  Remaining lifetime: 1130 secs, Level: 2, Interface: 101
  Estimated free bytes: 208, Actual free bytes: 208
  Aging timer expires in: 1130 secs

```

```

Packet: LSP ID: Router-E.02-00, Length: 76 bytes, Lifetime : 1161 secs
  Checksum: 0xb4fa, Sequence: 0x1, Attributes: 0x3 <L1 L2>
  NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
  Packet type: 20, Packet version: 1, Max area: 0

```

TLVs:

IS neighbor: Router-E.00, Internal, Metric: default 0
 IS neighbor: Router-A.00, Internal, Metric: default 0
 IS extended neighbor: Router-E.00, Metric: default 0
 IS extended neighbor: Router-A.00, Metric: default 0

No queued transmissions

Router-F.00-00 Sequence: 0x5, Checksum: 0x94bd, Lifetime: 1153 secs
 IS neighbor: Router-D.02 Metric: 10
 Two-way fragment: Router-D.02-00, Two-way first fragment: Router-D.02-00
 IS neighbor: Router-F.02 Metric: 10
 Two-way fragment: Router-F.02-00, Two-way first fragment: Router-F.02-00
 IP prefix: 10.0.0.16/30 Metric: 10 Internal Up
 IP prefix: 10.0.0.20/30 Metric: 10 Internal Up
 IP prefix: 192.168.0.6/32 Metric: 0 Internal Up

Header: LSP ID: Router-F.00-00, Length: 208 bytes
 Allocated length: 284 bytes, Router ID: 192.168.0.6
 Remaining lifetime: 1153 secs, Level: 2, Interface: 101
 Estimated free bytes: 76, Actual free bytes: 76
 Aging timer expires in: 1153 secs
 Protocols: IP, IPv6

Packet: LSP ID: Router-F.00-00, Length: 208 bytes, Lifetime : 1183 secs
 Checksum: 0x94bd, Sequence: 0x5, Attributes: 0x3 <L1 L2>
 NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
 Packet type: 20, Packet version: 1, Max area: 0

TLVs:

Area address: 49.0002 (3)
 LSP Buffer Size: 1492
 Speaks: IP
 Speaks: IPV6
 IP router id: 192.168.0.6
 IP address: 192.168.0.6
 Hostname: Router-F
 IP prefix: 192.168.0.6/32, Internal, Metric: default 0, Up
 IP prefix: 10.0.0.16/30, Internal, Metric: default 10, Up
 IP prefix: 10.0.0.20/30, Internal, Metric: default 10, Up
 IP extended prefix: 192.168.0.6/32 metric 0 up
 IP extended prefix: 10.0.0.16/30 metric 10 up
 IP extended prefix: 10.0.0.20/30 metric 10 up
 IS neighbor: Router-D.02, Internal, Metric: default 10
 IS neighbor: Router-F.02, Internal, Metric: default 10
 IS extended neighbor: Router-D.02, Metric: default 10
 IP address: 10.0.0.21
 Local interface index: 94, Remote interface index: 0
 IS extended neighbor: Router-F.02, Metric: default 10
 IP address: 10.0.0.18
 Local interface index: 93, Remote interface index: 0

No queued transmissions

Router-F.02-00 Sequence: 0x1, Checksum: 0xf5ae, Lifetime: 1153 secs
 IS neighbor: Router-E.00 Metric: 0
 Two-way fragment: Router-E.00-00, Two-way first fragment: Router-E.00-00
 IS neighbor: Router-F.00 Metric: 0
 Two-way fragment: Router-F.00-00, Two-way first fragment: Router-F.00-00

Header: LSP ID: Router-F.02-00, Length: 76 bytes
 Allocated length: 284 bytes, Router ID: 0.0.0.0
 Remaining lifetime: 1153 secs, Level: 2, Interface: 101

Estimated free bytes: 208, Actual free bytes: 208
Aging timer expires in: 1153 secs

Packet: LSP ID: Router-F.02-00, Length: 76 bytes, Lifetime : 1183 secs
Checksum: 0xf5ae, Sequence: 0x1, Attributes: 0x3 <L1 L2>
NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
Packet type: 20, Packet version: 1, Max area: 0

TLVs:

IS neighbor: Router-F.00, Internal, Metric: default 0
IS neighbor: Router-E.00, Internal, Metric: default 0
IS extended neighbor: Router-F.00, Metric: default 0
IS extended neighbor: Router-E.00, Metric: default 0
No queued transmissions

show isis hostname

Syntax	show isis hostname <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and QFX Series)	show isis hostname
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	<p>Display IS-IS hostname database information.</p> <p>This command displays the system ID-to-name cache. The output shows if the mapping has been learned by receipt of a Hostname TLV #137 (type dynamic) configured in Junos OS with the set system host-name command, or a static mapping defined in Junos OS with the set system static-host-mapping hostname sysid command (type static). The local router always has its type set to static even if static-host-mapping is not configured.</p>
Options	<p>none—Display IS-IS hostname database information.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show isis hostname on page 635
Output Fields	Table 89 on page 634 describes the output fields for the show isis hostname command. Output fields are listed in the approximate order in which they appear.

Table 89: show isis hostname Output Fields

Field Name	Field Description
System Id	System identifier mapped to the hostname.
Hostname	Hostname mapped to the system identifier.
Type	Type of mapping between system identifier and hostname. <ul style="list-style-type: none">• Dynamic—Hostname mapping determined as described in RFC 2763, <i>Dynamic Hostname Exchange Mechanism for IS-IS</i>.• Static—Hostname mapping configured by user.

Sample Output

show isis hostname

```
user@host> show isis hostname
IS-IS hostname database:
System Id      Hostname
1921.6800.4201 isis1
1921.6800.4202 isis2
1921.6800.4203 isis3
```

	Type
1921.6800.4201 isis1	Dynamic
1921.6800.4202 isis2	Static
1921.6800.4203 isis3	Dynamic

show isis interface


Syntax	show isis interface <brief detail extensive> <interface-name> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and QFX Series)	show isis interface <brief detail extensive> <interface-name>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display status information about IS-IS-enabled interfaces.
<div> NOTE: If the configured metric for an IS-IS level is above 63, and the wide-metrics-only statement is not configured, the show isis interface detail command and the show isis interface extensive command display 63 as the metric value for that level. Configure the wide-metrics-only statement to generate metric values greater than 63 on a per IS-IS level basis.</div> <p>The show isis interface command displays the configured metric value for an IS-IS level irrespective of whether is configured or not.</p>	
Options	none —Display standard information about all IS-IS-enabled interfaces. brief detail extensive —(Optional) Display the specified level of output. interface-name —(Optional) Display information about the specified interface only. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• <i>Example: Enabling Wide IS-IS Metrics for Traffic Engineering</i>
List of Sample Output	show isis interface on page 638 show isis interface brief on page 638 show isis interface detail on page 639 show isis interface extensive on page 639
Output Fields	Table 90 on page 637 describes the output fields for the show isis interface command. Output fields are listed in the approximate order in which they appear.

Table 90: show isis interface Output Fields

Field Name	Field Description	Level of Output
<i>interface-name</i>	Name of the interface.	detail
Designated router	Routing device selected by other routers that is responsible for sending link-state advertisements that describe the network. Used only on broadcast networks.	detail
Index	Interface index assigned by the Junos OS kernel.	detail
State	Internal implementation information.	detail
Circuit id	Circuit identifier.	detail
Circuit type	Circuit type: <ul style="list-style-type: none"> • 1—Level 1 only • 2—Level 2 only • 3—Level 1 and Level 2 	detail
LSP interval	Interval between link-state PDUs sent from the interface.	detail
CSNP interval	Interval between complete sequence number PDUs sent from the interface.	detail extensive
Sysid	System identifier.	detail
Interface	Interface through which the adjacency is made.	none brief
L or Level	Level: <ul style="list-style-type: none"> • 1—Level 1 only • 2—Level 2 only • 3—Level 1 and Level 2 	All levels
CirID	Circuit identifier.	none brief
Level 1 DR	Level 1 designated intermediate system.	none brief
Level 2 DR	Level 2 designated intermediate system.	none brief
L1/L2 Metric	Interface's metric for Level 1 and Level 2. If there is no information, the metric is 0.	none brief
Adjacency advertisement: Advertise	This routing device has signaled to advertise this interface to its neighbors in their label-switched paths (LSPs).	detail extensive
Adjacency advertisement: Suppress	This neighbor has signaled not to advertise this interface in the routing device's outbound LSPs.	detail extensive
Adjacencies	Number of adjacencies established on this interface.	detail

Table 90: show isis interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
Priority	Priority value for this interface.	detail
Metric	Metric value for this interface.	detail
Hello(s) / Hello Interval	Interface's hello interval.	detail extensive
Hold(s) / Hold Time	Interface's hold time.	detail extensive
Designated Router	Router responsible for sending network link-state advertisements, which describe all the routing devices attached to the network.	detail
Hello padding	Type of hello padding: <ul style="list-style-type: none"> • Adaptive—On point-to-point connections, the hello packets are padded from the initial detection of a new neighbor until the neighbor verifies the adjacency as Up in the adjacency state TLV. If the neighbor does not support the adjacency state TLV, then padding continues. On LAN connections, padding starts from the initial detection of a new neighbor until there is at least one active adjacency on the interface. • Loose—(Default) The hello packet is padded from the initial detection of a new neighbor until the adjacency transitions to the Up state. • Strict—Padding is performed on all interface types and for all adjacency states, and is continuous. 	extensive
LDP sync state	Current LDP synchronization state: in sync , in holddown , or not supported .	extensive
reason	Reason for being in the LDP sync state.	extensive
config holdtime	Configured value of the hold timer.	extensive
remaining	If the state is not in sync and the hold time is not infinity, then this field displays the remaining hold time in seconds.	extensive

Sample Output

show isis interface

```

user@host> show isis interface
IS-IS interface database:
Interface          L CirID Level 1 DR      Level 2 DR      L1/L2 Metric
at-2/3/0.0         3   0x1 Point to Point    Point to Point    10/10
1o0.0              0   0x1 Passive           Passive           0/0

```

show isis interface brief

The output for the **show isis interface brief** command is identical to that for the **show isis interface** command. For sample output, see [show isis interface on page 638](#).

show isis interface detail

```

user@host> show isis interface detail
IS-IS interface database:
at-2/3/0.0
  Index: 66, State: 0x6, Circuit id: 0x1, Circuit type: 3
  LSP interval: 100 ms, CSNP interval: 5 s
  Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
    1             1      64     10     9.000     27
    2             1      64     10     9.000     27
lo0.0
  Index: 64, State: 0x6, Circuit id: 0x1, Circuit type: 0
  LSP interval: 100 ms, CSNP interval: disabled
  Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
    1             0      64      0     0 Passive
    2             0      64      0     0 Passive

```

show isis interface extensive

```

user@host> show isis interface extensive
IS-IS interface database:
xe-6/1/0.0
  Index: 75, State: 0x6, Circuit id: 0x1, Circuit type: 2
  LSP interval: 100 ms, CSNP interval: 10 s, Loose Hello padding
  Adjacency advertisement: Advertise
  Level 1
    Adjacencies: 0, Priority: 64, Metric: 10
    Disabled
  Level 2
    Adjacencies: 1, Priority: 64, Metric: 10
    Hello Interval: 20.000 s, Hold Time: 60 s
    Designated Router: nemean.03

```

show isis overview

Syntax	show isis overview <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and QFX Series)	show isis overview <instance <i>instance-name</i> >
Release Information	Command introduced in Junos OS Release 8.5. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display IS-IS overview information.
Options	<p>none—Display standard overview information about IS-IS for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Display overview information for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show isis overview on page 642
Output Fields	Table 91 on page 640 lists the output fields for the show isis overview command. Output fields are listed in the approximate order in which they appear.

Table 91: show isis overview Output Fields

Field Name	Field Description
Instance	IS-IS routing instance.
Router ID	Router ID of the routing device.
Adjacency holddown	Adjacency holddown capability: enabled or disabled .
Maximum Areas	Maximum number of IS-IS areas advertised by the routing device.
LSP life time	Lifetime of the link-state PDU, in seconds.
Attached bit evaluation	Attached bit capability: enabled or disabled .
SPF delay	Delay before performing consecutive shortest-path-first (SPF) calculations.
SPF holddown	Delay before performing additional SPF calculations after the maximum number of consecutive SPF calculations is reached.

Table 91: show isis overview Output Fields (*continued*)

Field Name	Field Description
SPF rapid runs	Maximum number of SPF calculations that can be performed in succession before the holddown timer begins.
Overload bit at startup is set	Overload bit capability is enabled.
Overload high metrics	Overload high metrics capability: enabled or disabled .
Overload timeout	Time period after which overload is reset and the time that remains before the timer is set to expire.
Traffic engineering	Traffic engineering capability: enabled or disabled .
Restart	Graceful restart capability: enabled or disabled .
Restart duration	Time period for complete reacquisition of IS-IS neighbors.
Helper mode	Graceful restart helper capability: enabled or disabled .
Level	IS-IS level: <ul style="list-style-type: none"> • 1—Level 1 information • 2—Level 2 information
IPv4 is enabled	IP Protocol version 4 capability is enabled.
IPv6 is enabled	IP Protocol version 6 capability is enabled.
CLNS is enabled	(J Series routers only) OSI CLNP capability is enabled.
Internal route preference	Preference value of internal routes.
External route preference	Preference value of external routes.
Prefix export limit	Number of prefixes allowed to be exported, as configured by the prefix-export-limit statement.
Prefix export count	Number of prefixes exported.
Wide area metrics are enabled	Wide area metrics capability is enabled.
Narrow metrics are enabled	Narrow metrics capability is enabled.

Sample Output

show isis overview

```
user@host> show isis overview
Instance: master
  Router ID: 10.255.107.183
  Adjacency holddown: disabled
  Maximum Areas: 3
  LSP life time: 1200
  Attached bit evaluation: enabled
  SPF delay: 200 msec, SPF holddown: 5000 msec, SPF rapid runs: 3
  IPv4 is enabled, IPv6 is enabled
  Traffic engineering: enabled
  Restart: Disabled
    Helper mode: Enabled
Level 1
  Internal route preference: 15
  External route preference: 160
  Wide metrics are enabled, Narrow metrics are enabled
Level 2
  Internal route preference: 18
  External route preference: 165
  Prefix export limit: 5, Prefix export count: 5
  Wide metrics are enabled
```

show isis route

Syntax	<pre>show isis route <destination> <inet inet6> <instance instance-name> <logical-system (all logical-system-name)> <topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show isis route <destination> <inet inet6> <instance instance-name> <topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 12.1 for the QFX Series.</p>
Description	Display the routes in the IS-IS routing table.
Options	<p>none—Display all routes in the IS-IS routing table for all supported address families for all routing instances.</p> <p>destination—(Optional) Destination address for the route.</p> <p>inet inet6—(Optional) Display inet (IPv4) or inet6 (IPv6) routes, respectively.</p> <p>instance instance-name—(Optional) Display routes for the specified routing instance only.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)—(Optional) Display routes for the specified topology only, or use unicast to display information, if available, for both IPv4 and IPv6 unicast topologies.</p>
Required Privilege Level	view
List of Sample Output	<p>show isis route logical-system on page 644</p> <p>show isis route (CLNS) on page 644</p> <p>show isis route on page 645</p>
Output Fields	<p>Table 92 on page 643 describes the output fields for the show isis route command. Output fields are listed in the approximate order in which they appear.</p>

Table 92: show isis route Output Fields

Field Name	Field Description
Current version	Number of the current version of the IS-IS routing table.

Table 92: show isis route Output Fields (*continued*)

Field Name	Field Description
L1	Version of Level 1 SPF that was run.
L2	Version of Level 2 SPF that was run.
Prefix	Destination of the route.
L	IS-IS level: <ul style="list-style-type: none"> • 1—Level 1 only • 2—Level 2 only • 3—Level 1 and Level 2
Version	Version of SPF that generated the route.
Metric	Metric value associated with the route.
Type	Metric type: int (internal) or ext (external).
Interface	Interface to the next hop.
Via	System identifier of the next hop, displayed as a name if possible.
ISO Routes	ISO routing table entries.
snpa	MAC address.

Sample Output

show isis route logical-system

```

user@host> show isis route logical-system ls1
IS-IS routing table          Current version: L1: 8 L2: 11
Prefix      L Version Metric Type Interface  Via
10.9.7.0/30  2      11    20 int  gr-0/2/0.0  h
10.9.201.1/32 2      11    60 int  gr-0/2/0.0  h
IPv6 Unicast IS-IS routing table          Current version: L1: 9 L2: 11
Prefix      L Version Metric Type Interface  Via
8009:3::a09:3200/126 2      11    20 int  gr-0/2/0.0  h

```

show isis route (CLNS)

```

user@host> show isis route
IS-IS routing table          Current version: L1: 10 L2: 8
IPv4/IPv6 Routes
Prefix      L Version Metric Type Interface  Via
0.0.0.0/0    1      10    10 int  fe-0/0/1.0  ISIS.0
ISO Routes
Prefix L   Version Metric Type Interface  Via  snpa
0/0    1      10    10 int  fe-0/0/1.0  isis.0 0:12:0:34:0:56
47.0005.80ff.f800.0000.0108.0001/104

```

```

1          10          0 int
47.0005.80ff.f800.0000.0108.0001.1921.6800.4001/152
1          10          10 int fe-0/0/1.0 isis.0 0:12:0:34:0:56
47.0005.80ff.f800.0000.0108.0001.1921.6800.4002/152
1          10          20 int fe-0/0/1.0 isis.0 0:12:0:34:0:56
47.0005.80ff.f800.0000.0108.0002/104
1          10          0 int
47.0005.80ff.f800.0000.0108.0002.1921.6800.4001/152
1          10          10 int fe-0/0/1.0 isis.0 0:12:0:34:0:56

```

show isis route

```
user@host> show isis route
```

```

IS-IS routing table          Current version: L1: 4 L2: 13
IPv4/IPv6 Routes
-----
Prefix                      L   Version  Metric Type Interface      NH   Via
10.255.71.52/32             2    13        10   int  ae0.0                   IPV4 camaro
10.255.71.238/32           2    13        20   int  so-6/0/0.0             IPV4 olympic
                               as0.0                   IPV4 glacier
10.255.71.239/32           2    13        20   int  so-6/0/0.0             IPV4 olympic
                               ae0.0                   IPV4 camaro
10.255.71.242/32           2    13        10   int  as0.0                   IPV4 glacier
10.255.71.243/32           2    13        10   int  so-6/0/0.0             IPV4 olympic
12.13.0.0/30                2    13        20   int  so-6/0/0.0             IPV4 olympic
12.15.0.0/30                2    13        20   int  so-6/0/0.0             IPV4 olympic
13.15.0.0/30                2    13        30   int  ae0.0                   IPV4 camaro
                               so-6/0/0.0             IPV4 olympic
                               as0.0                   IPV4 glacier
13.16.0.0/30                2    13        25   int  as0.0                   IPV4 glacier
14.15.0.0/30                2    13        20   int  ae0.0                   IPV4 camaro
192.2.1.0/30                2    13        30   int  so-6/0/0.0             IPV4 olympic
                               as0.0                   IPV4 glacier
1eee::/64                   2    13        30   int  so-6/0/0.0             IPV6 olympic
                               as0.0                   IPV6 glacier
abcd::10:255:71:52/128     2    13        10   int  ae0.0                   IPV6 camaro
abcd::10:255:71:238/128    2    13        20   int  so-6/0/0.0             IPV6 olympic
                               as0.0                   IPV6 glacier
abcd::10:255:71:239/128    2    13        20   int  so-6/0/0.0             IPV6 olympic

```

					ae0.0	IPV6 camaro
abcd::10:255:71:242/128	2	13	10	int	as0.0	IPV6 glacier
abcd::10:255:71:243/128	2	13	10	int	so-6/0/0.0	IPV6 olympic

show isis spf

Syntax	show isis spf (brief log results) <instance <i>instance-name</i> > <level (1 2)> <logical-system (all <i>logical-system-name</i>)> <topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)>
Syntax (EX Series Switches)	show isis spf (brief log results) <instance <i>instance-name</i> > <level (1 2)> <topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display information about IS-IS shortest-path-first (SPF) calculations.
Options	<p>brief—Display an overview of SPF calculations.</p> <p>instance <i>instance instance-name</i>—(Optional) Display SPF calculations for the specified routing instance.</p> <p>level (1 2)—(Optional) Display SPF calculations for the specified IS-IS level.</p> <p>log—Display the log of SPF calculations.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>results—Display the results of SPF calculations.</p> <p>topology (ipv4-multicast ipv6-multicast ipv6-unicast unicast)—(Optional) Display SPF calculations for the specified topology only.</p>
Required Privilege Level	view
List of Sample Output	show isis spf log on page 648 show isis spf results logical-system on page 649 show isis spf results (CLNS) on page 650
Output Fields	Table 93 on page 647 describes the output fields for the show isis spf command. Output fields are listed in the approximate order in which they appear.

Table 93: show isis spf Output Fields

Field Name	Field Description
Node	System ID of a node.
Metric	Metric to the node.

Table 93: show isis spf Output Fields (*continued*)

Field Name	Field Description
Interface	Interface of the next hop.
Via	System ID of the next hop.
SNPA	Subnetwork point of attachment (MAC address of the next hop).
Start time	(log option only) Time that the SPF computation started.
Elapsed (secs)	(log option only) Length of time, in seconds, required to complete the SPF computation.
Count	(log option only) Number of times the SPF was triggered.
Reason	(log option only) Reason that the SPF computation was completed.

Sample Output

show isis spf log

```

user@host> show isis spf log logical-system lsl
IS-IS level 1 SPF log:
Start time           Elapsed (secs) Count Reason
Fri Oct 31 12:41:18   0.000069    1 Reconfig
Fri Oct 31 12:41:18   0.000107    3 Updated LSP fix.00-00
Fri Oct 31 12:41:18   0.000050    3 Address change on so-1/2/2.0
Fri Oct 31 12:41:23   0.000033    1 Updated LSP fix.00-00
Fri Oct 31 12:41:28   0.000178    5 New adjacency scat on ge-1/1/0.0
Fri Oct 31 12:41:59   0.000060    1 Updated LSP fix.00-00
Fri Oct 31 12:42:30   0.000161    2 Multi area attachment change
Fri Oct 31 12:56:58   0.000198    1 Periodic SPF
Fri Oct 31 13:10:29   0.000209    1 Periodic SPF
IS-IS level 2 SPF log:
Start time           Elapsed (secs) Count Reason
Fri Oct 31 12:41:18   0.000035    1 Reconfig
Fri Oct 31 12:41:18   0.000047    2 Updated LSP fix.00-00
Fri Oct 31 12:41:18   0.000043    5 Address change on gr-0/2/0.0
Fri Oct 31 12:41:23   0.000022    1 Updated LSP fix.00-00
Fri Oct 31 12:41:59   0.000144    3 New adjacency h on gr-0/2/0.0
Fri Oct 31 12:42:30   0.000257    3 New LSP skag.00-00
Fri Oct 31 12:54:37   0.000195    1 Periodic SPF
Fri Oct 31 12:55:50   0.000178    1 Updated LSP fix.00-00
Fri Oct 31 12:55:55   0.000174    1 Updated LSP h.00-00
Fri Oct 31 12:55:58   0.000176    1 Updated LSP skag.00-00
Fri Oct 31 13:08:14   0.000198    1 Periodic SPF
IPv6 Unicast IS-IS level 1 SPF log:
Start time           Elapsed (secs) Count Reason
Fri Oct 31 12:41:18   0.000028    1 Reconfig
Fri Oct 31 12:41:18   0.000043    3 Updated LSP fix.00-00
Fri Oct 31 12:41:18   0.000112    4 Updated LSP fix.00-00
Fri Oct 31 12:41:23   0.000059    1 Updated LSP fix.00-00
Fri Oct 31 12:41:25   0.000041    1 Updated LSP fix.00-00

```

```

Fri Oct 31 12:41:28      0.000103    5 New adjacency scat on ge-1/1/0.0
Fri Oct 31 12:41:59      0.000040    1 Updated LSP fix.00-00
Fri Oct 31 12:42:30      0.000118    2 Multi area attachment change
Fri Oct 31 12:56:08      0.000289    1 Periodic SPF
Fri Oct 31 13:11:07      0.000214    1 Periodic SPF
IPv6 Unicast IS-IS level 2 SPF log:

```

```

Start time      Elapsed (secs) Count Reason
Fri Oct 31 12:41:18      0.000027    1 Reconfig
Fri Oct 31 12:41:18      0.000039    2 Updated LSP fix.00-00
Fri Oct 31 12:41:18      0.000049    6 Updated LSP fix.00-00
Fri Oct 31 12:41:23      0.000025    1 Updated LSP fix.00-00
Fri Oct 31 12:41:25      0.000023    1 Updated LSP fix.00-00
Fri Oct 31 12:41:59      0.000087    3 New adjacency h on gr-0/2/0.0
Fri Oct 31 12:42:30      0.000123    3 New LSP skag.00-00
Fri Oct 31 12:55:50      0.000121    1 Updated LSP fix.00-00
Fri Oct 31 12:55:55      0.000121    1 Updated LSP h.00-00
Fri Oct 31 12:55:58      0.000121    1 Updated LSP skag.00-00
Fri Oct 31 13:09:46      0.000201    1 Periodic SPF
...

```

show isis spf results logical-system

```
user@host> show isis spf results logical-system ls1
```

IS-IS level 1 SPF results:

Node	Metric	Interface	Via	SNPA
scat.00	10	ge-1/1/0.0	scat	0:90:69:a6:48:9d
	20	10.9.1.0/30		
fix.02	10			
fix.00	0			
	10	10.9.1.0/30		
	10	10.9.5.0/30		
	10	10.9.6.0/30		
	20	10.9.7.0/30		
	60	10.9.201.1/32		

3 nodes

IS-IS level 2 SPF results:

Node	Metric	Interface	Via	SNPA
skag.00	20	gr-0/2/0.0	h	
	30	10.9.7.0/30		
skag.02	20	gr-0/2/0.0	h	
h.00	10	gr-0/2/0.0	h	
	20	10.9.6.0/30		
	20	10.9.7.0/30		
	60	10.9.201.1/32		
fix.00	0			
	10	10.9.1.0/30		
	10	10.9.5.0/30		
	10	10.9.6.0/30		

4 nodes

IPv6 Unicast IS-IS level 1 SPF results:

Node	Metric	Interface	Via	SNPA
scat.00	10	ge-1/1/0.0	scat	0:90:69:a6:48:9d
		ge-1/1/0.0	scat	0:90:69:a6:48:9d
	20	8009:1::a09:1400/126		
fix.02	10			
fix.00	0			
	10	8009:1::a09:1400/126		
	10	8009:2::a09:1e00/126		

```

                20      8009:3::a09:3200/126
                10      8009:4::a09:2800/126
    3 nodes

IPv6 Unicast IS-IS level 2 SPF results:
Node      Metric      Interface      Via      SNPA
skag.00    20      gr-0/2/0.0    h
           30      8009:3::a09:3200/126
skag.02    20      gr-0/2/0.0    h
           10      gr-0/2/0.0    h
h.00       10      gr-0/2/0.0    h
           20      8009:3::a09:3200/126
           20      8009:4::a09:2800/126
fix.00     0
           10      8009:1::a09:1400/126
           10      8009:2::a09:1e00/126
           10      8009:4::a09:2800/126
    4 nodes

Multicast IS-IS level 1 SPF results:
Node      Metric      Interface      Via      SNPA
scat.00    10      ge-1/1/0.0    scat    0:90:69:a6:48:9d
fix.02     10
fix.00     0
    3 nodes

Multicast IS-IS level 2 SPF results:
Node      Metric      Interface      Via      SNPA
skag.00    20      gr-0/2/0.0    h
skag.02    20      gr-0/2/0.0    h
h.00       10      gr-0/2/0.0    h
fix.00     0
    4 nodes
...

```

show isis spf results (CLNS)

```

user@host> show isis spf results
IS-IS level 1 SPF results:
Node      Metric      Interface      Via      SNPA
skag.00 10      fe-0/0/1.0    toothache 0:12:0:34:0:56
           10      fe-0/0/1.0    toothache 0:12:0:34:0:56
           20      192.168.37.64/29
           10      1921.6800.4001
           20      1921.6800.4002
pro1-a.02 10
pro1-a.00 0
           0      10.255.245.1/32
           10      192.168.37.64/29
           0      1921.6800.4211
    3 nodes

IS-IS level 2 SPF results:
Node      Metric      Interface      Via      SNPA
skag.00 10      fe-0/0/1.0    toothache 0:12:0:34:0:56
           10      fe-0/0/1.0    toothache 0:12:0:34:0:56
           20      10.255.245.1/32
           20      192.168.37.64/29
           20      47.0005.80ff.f800.0000.0109.0010/104

```

pro1-a.02	10	
pro1-a.00	0	
	0	10.255.245.1/32
	10	192.168.37.64/29
3 nodes		

show isis statistics

Syntax	show isis statistics <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and QFX Series)	show isis statistics <instance <i>instance-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display statistics about IS-IS traffic.
Options	none —Display IS-IS traffic statistics for all routing instances. instance <i>instance-name</i> —(Optional) Display statistics for the specified routing instance. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• clear isis statistics on page 604
List of Sample Output	show isis statistics on page 654
Output Fields	Table 94 on page 653 describes the output fields for the show isis statistics command. Output fields are listed in the approximate order in which they appear.

Table 94: show isis statistics Output Fields

Field Name	Field Description
PDU type	<p>PDU type:</p> <ul style="list-style-type: none"> • CSNP—Complete sequence number PDUs contain a complete list of all link-state PDUs in the IS-IS database. CSNPs are sent periodically on all links, and the receiving systems use the information in the CSNP to update and synchronize their link-state PDU databases. The designated router multicasts CSNPs on broadcast links in place of sending explicit acknowledgments for each link-state PDU. • IIH—IS-IS hello packets are broadcast to discover the identity of neighboring IS-IS systems and to determine whether the neighbors are Level 1 or Level 2 intermediate systems. • LSP—Link-state PDUs contain information about the state of adjacencies to neighboring IS-IS systems. Link-state PDUs are flooded periodically throughout an area. • PSNP—Partial sequence number PDUs are sent multicast by a receiver when it detects that it is missing a link-state PDU (when its link-state PDU database is out of date). The receiver sends a PSNP to the system that transmitted the CSNP, effectively requesting that the missing link-state PDU be transmitted. That routing device, in turn, forwards the missing link-state PDU to the requesting routing device. • Unknown—The PDU type is unknown.
Received	Number of PDUs received since IS-IS started or since the statistics were set to zero.
Processed	Number of PDUs received less the number dropped.
Drops	Number of PDUs dropped.
Sent	Number of PDUs transmitted since IS-IS started or since the statistics were set to zero.
Rexmit	Number of PDUs retransmitted since IS-IS started or since the statistics were set to zero.
Total packets received/sent	Total number of PDUs received and transmitted since IS-IS started or since the statistics were set to zero.
SNP queue length	Number of CSPN and PSNP packets currently waiting in the queue for processing. This value is almost always 0.
LSP queue length	Number of link-state PDUs waiting in the queue for processing. This value is almost always 0.
SPF runs	Number of shortest-path-first (SPF) calculations that have been performed. If this number is incrementing rapidly, it indicates that the network is unstable.
Fragments rebuilt	Number of link-state PDU fragments that the local system has computed.
LSP regenerations	Number of link-state PDUs that have been regenerated. A link-state PDU is regenerated when it is nearing the end of its lifetime and it has not changed.
Purges initiated	Number of purges that the system initiated. A purge is initiated if the software decides that a link-state PDU must be removed from the network.

Sample Output

show isis statistics

```
user@host> show isis statistics
```

```
IS-IS statistics for merino:
```

PDU type	Received	Processed	Drops	Sent	Rexmit
LSP	12227	12227	0	8184	683
IIH	113808	113808	0	115817	0
CSNP	198868	198868	0	198934	0
PSNP	6985	6979	6	8274	0
Unknown	0	0	0	0	0
Totals	331888	331882	6	331209	683

```
Total packets received: 331888 Sent: 331892
```

```
SNP queue length:      0 Drops:      0  
LSP queue length:      0 Drops:      0
```

```
SPF runs:              1014  
Fragments rebuilt:     1038  
LSP regenerations:     425  
Purges initiated:      0
```


CHAPTER 13

LLDP Operational Mode Commands

- `clear lldp neighbor`
- `clear lldp statistics`
- `show lldp`
- `show lldp local-information`
- `show lldp neighbors`
- `show lldp remote-global-statistics`
- `show lldp statistics`

clear lldp neighbor

Syntax	clear lldp neighbor <interface <i>interface-name</i>>
Release Information	Command introduced in Junos OS Release 9.6.
Description	<p>On MX Series and T Series routers, clear information regarding all Link Layer Discovery Protocol (LLDP) neighbors or LLDP neighbors of the specified interface.</p> <p>For information about interface names, see <i>Interface Naming Overview</i>. For information about interface names for TX Matrix routers, see <i>TX Matrix Router Chassis and Interface Names</i>. For information about FPC numbering on TX Matrix routers, see <i>Routing Matrix with a TX Matrix Router FPC Numbering</i>.</p>
Options	interface <i>interface-name</i> —(Optional) Clear the LLDP neighbors on the specified interface.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• clear lldp statistics on page 657
List of Sample Output	clear lldp statistics on page 656
Output Fields	When you enter this command, you are provided no feedback on the status of your request. You can enter the show lldp neighbors command before and after clearing the LLDP neighbors to verify the clear operation.

Sample Output

clear lldp statistics

```
user@host> clear lldp statistics
user@host> clear lldp statistics interface ge-0/2/0
```

clear lldp statistics

Syntax	clear lldpp neighbor <interface <i>interface-name</i>>
Release Information	Command introduced in Junos OS Release 9.6.
Description	<p>On MX Series and T Series routers, clear all Link Layer Discovery Protocols (LLDP) statistics or LLDP statistics associated with the specified interface.</p> <p>For information about interface names, see <i>Interface Naming Overview</i>. For information about interface names for TX Matrix routers, see <i>TX Matrix Router Chassis and Interface Names</i>. For information about FPC numbering on TX Matrix routers, see <i>Routing Matrix with a TX Matrix Router FPC Numbering</i>.</p>
Options	interface <i>interface-name</i> —(Optional) Clear LLDP statistics on the specified interface.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • clear lldp neighbor on page 656
List of Sample Output	clear lldp neighbor on page 657
Output Fields	When you enter this command, you are provided no feedback on the status of your request. You can enter the show lldp statistics command before and after clearing the LLDP statistics to verify the clear operation.

Sample Output

clear lldp neighbor

```
user@host> clear lldp neighbors
user@host> clear lldp neighbors interface ge-0/2/2
```

show lldp

Syntax	show lldp <detail>
Release Information	Command introduced in Junos OS Release 9.6.
Description	On MX Series and T Series routers, display information about the Link Layer Discovery Protocol (LLDP).
Options	detail —(Optional) Display the detailed output level.
Required Privilege Level	view
List of Sample Output	show lldp on page 660 show lldp detail on page 660
Output Fields	Table 95 on page 658 describes the output fields for the show lldp command. Output fields are listed in the approximate order in which they appear.

Table 95: show lldp Output Fields

Field Name	Field Description
LLDP	Status of LLDP: Enabled or Disabled .
Advertisement interval	Value of the advertisement interval parameter.
Transmit delay	Value of the transmit delay parameter.
Hold timer	Value of the hold timer parameter.
Notification interval	Value of the notification interval parameter.
Config Trap Interval	Value of the configuration trap parameter.
Connection Hold timer	Value of the connection hold timer parameter.
Interface	<p>List of LLDP interfaces, showing status (Enabled or Disabled) and Neighbor count (detail only).</p> <p>For information about interface names, see <i>Interface Naming Overview</i>. For information about interface names for TX Matrix routers, see <i>TX Matrix Router Chassis and Interface Names</i>. For information about FPC numbering on TX Matrix routers, see <i>Routing Matrix with a TX Matrix Router FPC Numbering</i>.</p>
LLDP basic TLVs supported	List of basic LLDP TLVs supported by this device (detail only).

Table 95: show lldp Output Fields (*continued*)

Field Name	Field Description
LLDP 802 TLVs supported	List of IEEE 802.1 LLDP TLVs supported by this device (detail only).

Sample Output

show lldp

```
user@host> show lldp
LLDP : Enabled
Advertisement interval : 30 Second(s)
Transmit delay : 2 Second(s)
Hold timer : 4 Second(s)
Notification interval : 30 Second(s)
Config Trap Interval : 300 Second(s)
Connection Hold timer : 60 Second(s)

Interface      LLDP
ge-0/0/0       Enabled
ge-0/0/1       Enabled
ge-0/0/4       Enabled
```

Sample Output

show lldp detail

```
user@host> show lldp detail
LLDP : Enabled
Advertisement interval : 30 Second(s)
Transmit delay : 2 Second(s)
Hold timer : 4 Second(s)
Notification interval : 30 Second(s)
Config Trap Interval : 300 Second(s)
Connection Hold timer : 60 Second(s)

Interface      LLDP      Neighbor count
ge-0/0/0       Enabled   0
ge-0/0/1       Enabled   0
ge-0/0/4       Enabled   0
```

LLDP basic TLVs supported:

Chassis identifier, Port identifier, Port description, System name, System description, System capabilities, Management address.

LLDP 802 TLVs supported:

Link aggregation, Maximum frame size, MAC/PHY Configuration/Status, Port VLAN ID, Port VLAN name.

show lldp local-information

Syntax	show lldp local-information
Release Information	Command introduced in Junos OS Release 9.6.
Description	On MX Series and T Series routers, display local Link Layer Discovery Protocol (LLDP) information.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show lldp local-information(Management Information Address Subtype is IPv4) on page 663 show lldp local-information(Management Information Address Subtype is IPv6) on page 663
Output Fields	Table 96 on page 661 describes the output fields for the show lldp local-information command. Output fields are listed in the approximate order in which they appear.

Table 96: show lldp local-information Output Fields

Field Name	Field Description
LLDP Local Information details	Information that follows pertains to the local system.
Chassis ID	List of chassis identifiers for local information.
System name	Local system name reported by LLDP.
System descr	Local system description reported by LLDP.
System Capabilities	Capabilities (such as Bridge or Router) that are Supported or Enabled by system on the interface.
Management Information	Listed by Interface Name , Address Subtype (such as ipv4 , ipv6), Address (such as 192.168.168.229 , 1fd::1a10), Interface Number , and Interface Numbering Subtype .
Interface Name	List of local interfaces. For information about interface names, see <i>Interface Naming Overview</i> . For information about interface names for TX Matrix routers, see <i>TX Matrix Router Chassis and Interface Names</i> . For information about FPC numbering on TX Matrix routers, see <i>Routing Matrix with a TX Matrix Router FPC Numbering</i> .
Parent Interface	Name of the ae interface to which the interface belongs
Interface ID	List of local interface identifiers.

Table 96: show lldp local-information Output Fields (*continued*)

Field Name	Field Description
Interface Description	List of local interface descriptions.
Status	List of interface conditions: UP or DOWN .

Sample Output

show lldp local-information(Management Information Address Subtype is IPv4)

```

user@host> show lldp local-information
LLDP Local Information details

Chassis ID   : 64:87:88:65:37:c0
System name  : apg-hp1
System descr : Juniper Networks, Inc. mx240 , version 14.1I20131231_0701_builder
[builder] Build date: 2013-12-31 07:13:42 UTC

System Capabilities
  Supported      : Bridge Router
  Enabled        : Bridge Router

Management Information
  Interface Name : Unknown
  Address Subtype : IPv4(1)
  Address        : 10.216.97.103
  Interface Number : 1
  Interface Numbering Subtype : ifIndex(2)

Interface name  Parent Interface  Interface ID  Interface description  Status
ge-2/0/0       ae0                1475         ge-2/0/0              Up
ge-2/0/1       ae0                1476         ge-2/0/1              Up

```

show lldp local-information(Management Information Address Subtype is IPv6)

```

user@host> show lldp local-information
LLDP Local Information details

Chassis ID   : ac:4b:c8:92:67:c0
System name  : apg-hp
System descr : Juniper Networks, Inc. mx240 , version 13.2-20131210.0 [builder]
Build date: 2013-12-10 06:23:15 UTC

System Capabilities
  Supported      : Bridge Router
  Enabled        : Bridge Router

Management Information
  Interface Name : fxp0
  Address Subtype : IPv6(2)
  Address        : 1fd::1a20
  Interface Number : 1
  Interface Numbering Subtype : ifIndex(2)

Interface name  Parent Interface  Interface ID  Interface description  Status
ge-1/2/4       -                530          -                    Down
ge-1/2/5       -                531          -                    Down
ge-1/2/2       -                528          ge-1/2/2            Up
ge-1/2/3       -                529          ge-1/2/3            Up

```

show lldp neighbors

Syntax	<code>show lldp neighbors</code> <code><interface <i>interface-name</i>></code>
Release Information	Command introduced in Junos OS Release 9.6.
Description	On MX Series and T Series routers, display information about LLDP neighbors. For information about interface names, see <i>Interface Naming Overview</i> . For information about interface names for TX Matrix routers, see <i>TX Matrix Router Chassis and Interface Names</i> . For information about FPC numbering on TX Matrix routers, see <i>Routing Matrix with a TX Matrix Router FPC Numbering</i> .
Options	<code>interface <i>interface-name</i></code> —(Optional) Display the neighbor information about a particular physical interface.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> clear lldp neighbor on page 656
List of Sample Output	show lldp neighbors on page 666 show lldp neighbors interface ge-0/0/4 (Management Address is IPv4) on page 666 show lldp neighbors interface ge-0/0/4 (Management Address is IPv6) on page 667
Output Fields	Table 97 on page 664 describes the output fields for the show lldp neighbors command. Output fields are listed in the approximate order in which they appear.

Table 97: show lldp neighbors Output Fields

Field Name	Field Description
LLDP Remote Devices Information	Information about remote devices.
LocalInterface	List of local interfaces for which neighbor information is available.
ChassisId	List of chassis identifiers for neighbors.
PortInfo	List of port information gathered from neighbors. This could be the port identifier or port description.
SysName	List of system names gathered from neighbors.
LLDP Neighbor Information	Information about both local and neighbor systems on the interface (appears when the interface option is used).
Local Information	Information about local systems on the interface (appears when the interface option is used).

Table 97: show lldp neighbors Output Fields (*continued*)

Field Name	Field Description
Neighbor Information	Information about both local and neighbor system on the interface (appears when the interface option is used).
Index	Local interface index (appears when the interface option is used).
Time Mark	Date and timestamp of information (appears when the interface option is used).
Time To Live	Number of seconds for which this information is valid (appears when the interface option is used).
Local Interface	Name of the local physical interface (appears when the interface option is used).
Parent Interface	Name of the ae interface to which the interface belongs
Local Port ID	Local port identifier (appears when the interface option is used).
Neighbor Information	Information about neighbor systems on the interface (appears when the interface option is used).
Chassis type	Type of chassis identifier supplied, such as MAC address (appears when the interface option is used).
Chassis ID	Chassis identifier of type listed (appears when the interface option is used).
Port type	Type of port identifier supplied, such as local (appears when the interface option is used).
Port ID	Port identifier of type listed (appears when the interface option is used).
Port description	Port description (appears when the interface option is used).
System name	Name supplied by the system on the interface (appears when the interface option is used).
System Description	Description supplied by the system on the interface (appears when the interface option is used).
System Capabilities	Capabilities (such as bridge or router) that are Supported or Enabled by the system on the interface (appears when the interface option is used).
Management address	Details of the management address: Address Type (such as ipv4 and ipv6), Address (such as 10.204.34.35 , 1fd::1a10), Interface Number , Interface Subtype , and Organization Identifier (OID) (appears when the interface option is used).
Organization Info	One or more entries listing remote information by Organizationally Unique Identifier (OUI), Subtype , Index , and Info (appears when the interface option is used).

Sample Output

show lldp neighbors

```
user@host> show lldp neighbors
```

Local Interface	Parent Interface	Chassis Id	Port info	System Name
ge-2/0/0	ae0	ac:4b:c8:92:67:c0	528	apg-hp
ge-2/0/1	ae0	ac:4b:c8:92:67:c0	529	apg-hp

Sample Output

show lldp neighbors interface ge-0/0/4 (Management Address is IPv4)

```
user@host> show lldp neighbors interface ge-0/0/4
```

LLDP Neighbor Information:

Local Information:

Index: 2 Time to live: 120 Time mark: Tue Dec 31 11:47:46 2013 Age: 15 secs

Local Interface : ge-2/0/1

Parent Interface : ae0

Local Port ID : 1476

Ageout Count : 0

Neighbour Information:

Chassis type : Mac address

Chassis ID : ac:4b:c8:92:67:c0

Port type : Locally assigned

Port ID : 529

Port description : ge-1/2/3

System name : apg-hp

System Description : Juniper Networks, Inc. mx240 , version 14.1-20131222.0

[builder] Build date: 2013-12-22 09:13:26 UTC

System capabilities

Supported: Bridge Router

Enabled : Bridge Router

Management address

Address Type : IPv4(1)

Address : 10.216.98.57

Interface Number : 1

Interface Subtype : ifIndex(2)

OID : 1.3.6.1.2.1.31.1.1.1.1.1.

Organization Info

OUI : IEEE 802.3 Private (0x00120f)

Subtype : MAC/PHY Configuration/Status (1)

Info : Autonegotiation [supported, enabled (0x3)], PMD Autonegotiation

Capability (0x1d), MAU Type (0x0)

Index : 1

Organization Info

OUI : IEEE 802.3 Private (0x00120f)

Subtype : Link Aggregation (3)

Info : Aggregation Status (0x3), Aggregation Port ID (1694498816)

Index : 2

Organization Info

OUI : IEEE 802.3 Private (0x00120f)

```

Subtype : Maximum Frame Size (4)
Info    : MTU Size (1518)
Index   : 3

```

show lldp neighbors interface ge-0/0/4 (Management Address is IPv6)

```

user@host> show lldp neighbors interface ge-0/0/4
LLDP Neighbor Information:
Local Information:
Index: 1 Time to live: 120 Time mark: Thu Dec 12 07:19:45 2013 Age: 28 secs
Local Interface   : ge-1/2/2
Parent Interface  : -
Local Port ID     : 528
Ageout Count      : 0

Neighbour Information:
Chassis type      : Mac address
Chassis ID        : 64:87:88:65:37:c0
Port type         : Locally assigned
Port ID           : 1475
Port description  : ge-2/0/0
System name       : apg-hp1

System Description : Juniper Networks, Inc. mx240 , version 11.4R10 Build date:
2013-10-24 10:10:02 UTC

System capabilities
Supported: Bridge Router
Enabled  : Bridge Router

Management address
Address Type      : IPv6(2)
Address           : 1fd::1a10
Interface Number  : 1
Interface Subtype : ifIndex(2)
OID               : 1.3.6.1.2.1.31.1.1.1.1.1.

Organization Info
OUI               : IEEE 802.3 Private (0x00120f)
Subtype           : MAC/PHY Configuration/Status (1)
Info              : Autonegotiation [supported, enabled (0x3)], PMD Autonegotiation
Capability (0x5), MAU Type (0x0)
Index             : 1

Organization Info
OUI               : IEEE 802.3 Private (0x00120f)
Subtype           : Link Aggregation (3)
Info              : Aggregation Status (0x1), Aggregation Port ID (0)
Index             : 2

Organization Info
OUI               : IEEE 802.3 Private (0x00120f)
Subtype           : Maximum Frame Size (4)
Info              : MTU Size (1518)
Index             : 3

Organization Info
OUI               : Ethernet Bridged (0x0080c2)
Subtype           : VLAN Name (3)

```

Info : VLAN ID (100), VLAN Name (vlan-100)
Index : 4

show lldp remote-global-statistics

Syntax	show lldp remote-global-statistics
Release Information	Command introduced in Junos OS Release 9.6.
Description	On MX Series and T Series routers, display remote Link Layer Discovery Protocol (LLDP) global statistics.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show lldp remote-global-statistics on page 670
Output Fields	Table 98 on page 669 describes the output fields for the show lldp remote-global-statistics command. Output fields are listed in the approximate order in which they appear.

Table 98: show lldp remote-global-statistics Output Fields

Field Name	Field Description
LLDP Remote Database Table Counters	Information about remote database table counters.
LastchangeTime	Time elapsed between LLDP agent startup and the last change to the remote database table information.
Inserts	Number of insertions made in the remote database table.
Deletes	Number of deletions made in the remote database table.
Drops	Number of LLDP frames dropped from the remote database table because of errors.
Ageouts	Number of remote database table entries that have aged out of the table.

Sample Output

show lldp remote-global-statistics

```
user@host> show lldp remote-global-statistics
user@host> show lldp remote-global-statistics
LLDP Remote Database Table Counters
LastchangeTime      Inserts    Deletes    Drops    Ageouts
00:00:76 (76 sec)   192        0          0        0
```


show lldp statistics

Syntax	show lldp statistics <interface <i>interface-name</i> >
Release Information	Command introduced in Junos OS Release 9.6.
Description	On MX Series and T Series routers, display information about Link Layer Discovery Protocol (LLDP) statistics.
Options	interface <i>interface-name</i> —(Optional) Display the statistics about a particular physical interface.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> clear lldp statistics on page 657
List of Sample Output	show lldp statistics on page 672 show lldp statistics interface ge-0/1/1 on page 672
Output Fields	Table 99 on page 671 describes the output fields for the show lldp statistics command. Output fields are listed in the approximate order in which they appear.

Table 99: show lldp statistics Output Fields

Field Name	Field Description
Interface	Interface name. For information about interface names, see <i>Interface Naming Overview</i> . For information about interface names for TX Matrix routers, see <i>TX Matrix Router Chassis and Interface Names</i> . For information about FPC numbering on TX Matrix routers, see <i>Routing Matrix with a TX Matrix Router FPC Numbering</i> .
Received	Number of LLDP frames received on this interface.
Transmitted	Number of LLDP frames sent on this interface.
Unknown-TLVs	Number of LLDP frames with unsupported content received on this interface.
With-Errors	Number of LLDP frames with errors received on this interface.
Discarded	Number of LLDP frames received on this interface that were discarded because of problems.

Sample Output

show lldp statistics

```
user@host> show lldp statistics
Interface Received Transmitted Unknown-TLVs With-Errors Discarded
-----
ge-0/1/1 544 540 0 0 0
ge-0/1/2 540 500 0 0 0
ge-0/1/3 544 540 0 0 0
ge-0/1/4 544 540 0 0 0
ge-0/1/5 544 540 0 0 0
ge-0/1/6 544 540 0 0 0
ge-0/1/7 0 0 0 0 0
```

Sample Output

show lldp statistics interface ge-0/1/1

```
user@host> show lldp statistics interface ge-0/1/1
Interface Received Transmitted Unknown-TLVs With-Errors Discarded
-----
ge-0/1/1 544 540 0 0 0
```

CHAPTER 14

MVRP Operational Mode Commands

- `show mvrp`
- `show mvrp applicant-state`
- `show mvrp dynamic-vlan-memberships`
- `show mvrp interface`
- `show mvrp registration-state`
- `show mvrp statistics`

show mvrp

Syntax	show mvrp
Release Information	Command introduced in Junos OS Release 10.1. Command introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
Description	Display Multiple VLAN Registration Protocol (MVRP) configuration information.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show mvrp applicant-state on page 676• show mvrp dynamic-vlan-memberships on page 678• show mvrp interface on page 679• show mvrp registration-state on page 680• show mvrp statistics on page 682
List of Sample Output	show mvrp on page 674
Output Fields	Table 100 on page 674 lists the output fields for the show mvrp command. Output fields are listed in the approximate order in which they appear.

Table 100: show mvrp Output Fields

Field Name	Field Description
MVRP dynamic VLAN creation	Displays whether global MVRP dynamic Virtual LAN (VLAN) creation is Enabled or Disabled .
MVRP BPDU MAC address	Displays the multicast media access control (MAC) address for MVRP. If configured, the provider MVRP multicast MAC address is used; otherwise, the customer MVRP multicast MAC address is used.
MVRP timers (ms)	Displays MVRP timer information: <ul style="list-style-type: none">• Interface—The interface on which MVRP is configured.• Join—The maximum number of milliseconds the interfaces must wait before sending VLAN advertisements.• Leave—The number of milliseconds an interface must wait after receiving a Leave message to remove the interface from the VLAN specified in the message.• LeaveAll—The interval at which LeaveAll messages are sent on interfaces. LeaveAll messages maintain current MVRP VLAN membership information in the network.

Sample Output

show mvrp

```
user@host> show mvrp
```

```
MVRP configuration for routing instance 'default-switch'
MVRP dynamic VLAN creation : Enabled
MVRP BPDU MAC address      : Customer bridge group (01-80-C2-00-00-21)
MVRP timers (ms)
  Interface      Join   Leave  LeaveAll
  ge-11/2/8      200    800    10000
  ge-11/0/9      200    800    10000
  ge-11/3/0      200    800    10000
```

show mvrp applicant-state

Syntax	show mvrp applicant-state
Release Information	Command introduced in Junos OS Release 10.1.
Description	For MX Series routers and EX Series switches, display Multiple VLAN Registration Protocol (MVRP) applicant state information.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show mvrp on page 674• show mvrp interface on page 679• show mvrp registration-state on page 680• show mvrp statistics on page 682
List of Sample Output	show mvrp applicant-state on page 677
Output Fields	Table 101 on page 676 lists the output fields for the show mvrp applicant-state command. Output fields are listed in the approximate order in which they appear.

Table 101: show mvrp applicant-state Output Fields

Field Name	Field Description
VLAN Id	Displays the Virtual LAN (VLAN) ID number.
Interface	Displays the interface number associated with the VLAN ID.
State	Displays one of the following MVRP registrar states: <ul style="list-style-type: none">• VO—Very anxious observer.• VP—Very anxious passive.• VA—Very anxious new.• AN—Anxious new.• AA—Anxious active.• QA—Quiet active.• LA—Leaving active.• AO—Anxious observer.• QO—Quiet observer.• LO—Leaving observer.• AP—Anxious passive.• QA—Quiet passive.

Sample Output

show mvrp applicant-state

```
user@host> show mvrp applicant-state
MVRP applicant state for routing instance 'default-switch'
(V0) Very anxious observer, (VP) Very anxious passive, (VA) Very anxious new,
(AN) Anxious new, (AA) Anxious active, (QA) Quiet active, (LA) Leaving active,
(AO) Anxious observer, (QO) Quiet observer, (LO) Leaving observer,
(AP) Anxious passive, (QP) Quiet passive

VLAN Id      Interface      State
-----
100          ge-11/3/0      Declaring (QA)
200          ge-11/3/0      Declaring (QA)
300          ge-11/3/0      Declaring (QA)
```

show mvrp dynamic-vlan-memberships

Syntax	show mvrp dynamic-vlan-memberships
Release Information	Command introduced in Junos OS Release 10.1. Command introduced in Junos OS Release 13.2X50-D10 for EX Series switches.
Description	Display all Virtual LANs (VLANs) that have been created dynamically using Multiple VLAN Registration Protocol (MVRP) on the router or switch.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show mvrp on page 674 • show mvrp applicant-state on page 676 • show mvrp interface on page 679 • show mvrp registration-state on page 680 • show mvrp statistics on page 682
List of Sample Output	show mvrp dynamic-vlan-memberships on page 678
Output Fields	Table 102 on page 678 lists the output fields for the show mvrp dynamic-vlan-memberships command. Output fields are listed in the approximate order in which they appear.

Table 102: show mvrp dynamic-vlan-memberships Output Fields

Field Name	Field Description
VLAN Id	The VLAN ID of the dynamically created VLAN.
Interfaces	The interface or interfaces that are bound to the dynamically created VLAN.

Sample Output

show mvrp dynamic-vlan-memberships

```

user@host> show mvrp dynamic-vlan-memberships
MVRP dynamic vlans for routing instance 'default-switch'
(s) static vlan, (f) fixed registration

VLAN Id      Interfaces
  100 (s)    ge-11/3/0
  200 (s)    ge-11/3/0
  300 (s)

```


show mvrp interface

Syntax	show mvrp interface
Release Information	Command introduced in Junos OS Release 10.1.
Description	For MX Series routers and EX Series switches, display Multiple VLAN Registration Protocol (MVRP) interface-specific information.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show mvrp on page 674 • show mvrp applicant-state on page 676 • show mvrp dynamic-vlan-memberships on page 678 • show mvrp registration-state on page 680 • show mvrp statistics on page 682
List of Sample Output	show mvrp interface on page 679
Output Fields	Table 103 on page 679 lists the output fields for the show mvrp interface command. Output fields are listed in the approximate order in which they appear.

Table 103: show mvrp interface Output Fields

Field Name	Field Description
Interface	Interface on which MVRP is configured.
Status	Status of the MVRP: Enabled or Disabled .
Registration Mode	Registration for the interface: Fixed , Forbidden , or Normal .
Applicant Mode	Applicant mode.

Sample Output

show mvrp interface

```

user@host> show mvrp interface
MVRP interface information for routing instance 'default-switch'

Interface      Status      Registration Mode      Applicant
Mode
ge-11/2/8      Enabled     Normal                Normal
ge-11/0/9      Enabled     Normal                Normal
ge-11/3/0      Enabled     Normal                Normal

```

show mvrp registration-state

Syntax	show mvrp registration-state
Release Information	Command introduced in Junos OS Release 10.1.
Description	For MX Series routers and EX Series switches, display Multiple VLAN Registration Protocol (MVRP) registration state information.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show mvrp on page 674 • show mvrp dynamic-vlan-memberships on page 678 • show mvrp interface on page 679 • show mvrp statistics on page 682
List of Sample Output	show mvrp registration-state on page 680
Output Fields	Table 104 on page 680 lists the output fields for the show mvrp registration-state command. Output fields are listed in the approximate order in which they appear.

Table 104: show mvrp registration-state Output Fields

Field Name	Field Description
VLAN Id	Displays the Virtual LAN (VLAN) ID number.
Interface	Displays the interface number associated with the VLAN ID.
Registrar State	Displays whether the registrar state is Registered or Empty .
Forced State	Displays whether the forced state is Registered or Empty .
Managed State	Displays one of the following states: <ul style="list-style-type: none"> • Fixed—VLANs always stay in a registered state and are declared as such on all other forwarding ports. • Normal—VLANs participate in the MVRP protocol and honor incoming join requests normally. • Forbidden—VLANs ignore the incoming join requests and always stay in an unregistered state.
STP State	Displays whether the Spanning Tree Protocol (STP) is Blocking or Forwarding .

Sample Output

show mvrp registration-state

```
user@host> show mvrp registration-state
MVRP registration state for routing instance 'default-switch'
```

VLAN Id	Interface	Registrar State	Forced State	Managed State	STP State
100	ge-11/2/8	Empty	Registered	Fixed	Forwarding
	ge-11/0/9	Empty	Empty	Normal	Forwarding
	ge-11/3/0	Registered	Registered	Normal	Forwarding
101	ge-11/2/8	Empty	Registered	Fixed	Forwarding
	ge-11/0/9	Empty	Empty	Normal	Forwarding
	ge-11/3/0	Registered	Registered	Normal	Forwarding

show mvrp statistics

Syntax	show mvrp statistics
Release Information	Command introduced in Junos OS Release 10.1.
Description	For MX Series routers and EX Series switches, display Multiple VLAN Registration Protocol (MVRP) statistics in the form of Multiple Registration Protocol data unit (MRPDU) messages.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show mvrp on page 674• show mvrp applicant-state on page 676• show mvrp dynamic-vlan-memberships on page 678• show mvrp interface on page 679• show mvrp registration-state on page 680
List of Sample Output	show mvrp statistics on page 682
Output Fields	Table 105 on page 682 lists the output fields for the show mvrp statistics command. Output fields are listed in the approximate order in which they appear.

Table 105: show mvrp statistics Output Fields

Field Name	Field Description
interface name	Interface for which MVRP statistics are displayed.
VLAN IDs registered	Number of Virtual LAN (VLAN) IDs registered.
Sent MVRP PDUs	Number of MRPDU messages transmitted from the router.
Received MVRP PDUs without error	Number of MRPDU messages received on the router.
Received MVRP PDUs with error	Number of invalid MRPDU messages received on the router.

Sample Output

show mvrp statistics

```
user@host> show mvrp statistics
MVRP statistics for routing instance 'default-switch'

Interface name           : ge-11/2/8
VLAN IDs registered      : 0
Sent MVRP PDUs           : 1467
```

Received MVRP PDUs without error: 0
Received MVRP PDUs with error : 0

Interface name : ge-11/0/9
VLAN IDs registered : 0
Sent MVRP PDUs : 1418
Received MVRP PDUs without error: 702
Received MVRP PDUs with error : 0

Interface name : ge-11/3/0
VLAN IDs registered : 2
Sent MVRP PDUs : 1524
Received MVRP PDUs without error: 1366
Received MVRP PDUs with error : 0

CHAPTER 15

OSPF Operational Mode Commands

- `clear (ospf | ospf3) database`
- `clear (ospf | ospf3) io-statistics`
- `clear (ospf | ospf3) neighbor`
- `clear (ospf | ospf3) overload`
- `clear (ospf | ospf3) statistics`
- `show (ospf | ospf3) backup coverage`
- `show (ospf | ospf3) backup lsp`
- `show (ospf | ospf3) backup neighbor`
- `show (ospf | ospf3) backup spf`
- `show ospf context-identifier`
- `show ospf database`
- `show ospf3 database`
- `show (ospf | ospf3) interface`
- `show (ospf | ospf3) io-statistics`
- `show (ospf | ospf3) log`
- `show (ospf | ospf3) neighbor`
- `show (ospf | ospf3) overview`
- `show (ospf | ospf3) route`
- `show (ospf | ospf3) statistics`

clear (ospf | ospf3) database

Syntax clear (ospf | ospf3) database
<advertising-router (*router-id* | self)>
<area *area-id*>
<asbrsummary>
<external>
<instance *instance-name*>
<inter-area-prefix>
<inter-area-router>
<intra-area-prefix>
<link-local>
<logical-system (all | *logical-system-name*)>
<lsa-id *lsa-id*>
<netsummary>
<network>
<nssa>
<opaque-area>
<purge>
<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>
<router>

Syntax (EX Series Switch and QFX Series) clear (ospf | ospf3) database
<advertising-router (*router-id* | self)>
<area *area-id*>
<asbrsummary>
<external>
<instance *instance-name*>
<inter-area-prefix>
<inter-area-router>
<intra-area-prefix>
<link-local>
<lsa-id *lsa-id*>
<netsummary>
<network>
<nssa>
<opaque-area>
<purge>
<router>

Release Information Command introduced before Junos OS Release 7.4.
advertising-router *router-id*, **netsummary**, **network**, **nssa**, **opaque-area**, and **router** options added in Junos OS Release 8.3. You must use the **purge** command with these options.
area *area-id* option added in Junos OS Release 8.3.
Command introduced in Junos OS Release 9.0 for EX Series switches.
realm option added in Junos OS Release 9.2.
advertising-router (*router-id* | **self**) option added in Junos OS Release 9.5.
advertising-router (*router-id* | **self**) option introduced in Junos OS Release 9.5 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
purge option (and all options that are dependent on the **purge** option) hidden in Junos OS Release 13.3.

Description With the master Routing Engine, delete entries in the Open Shortest Path First (OSPF) link-state advertisement (LSA) database. With the backup Routing Engine, delete the OSPF LSA database and sync the new database with the master Routing Engine.



CAUTION: You can also use the **purge** command with any of the options to discard rather than delete the specified LSA entries. This command is useful only for testing. Use it with care, because it causes significant network disruption.

Options

- none**—Delete all LSAs other than the system's own LSAs, which are regenerated. To resynchronize the database, the system destroys all adjacent neighbors that are in the state **EXSTART** or higher. The neighbors are then reacquired and the databases are synchronized.
- advertising-router (*router-id* | *self*)**—(Hidden) Discard entries for the LSA entries advertised by the specified routing device or by this routing device.
- area *area-id***—(Optional) Discard entries for the LSAs in the specified area.
- asbrsummary**—(Optional) Discard summary AS boundary router LSA entries.
- external**—(Optional) Discard external LSAs.
- instance *instance-name***—(Optional) Delete or discard entries for the specified routing instance only.
- inter-area-prefix**—(OSPFv3 only) (Optional) Discard interarea prefix LSAs.
- inter-area-router**—(OSPFv3 only) (Optional) Discard interarea router LSAs.
- intra-area-prefix**—(OSPFv3 only) (Optional) Discard intra-area prefix LSAs.
- logical-system (*all* | *logical-system-name*)**—(Optional) Perform this operation on all logical systems or on a particular logical system.
- link-local**—(Optional) Delete link-local LSAs.
- lsa-id *lsa-id***—(Optional) Discard the LSA entries with the specified LSA identifier.
- netsummary**—(Hidden) Discard summary network LSAs.
- network**—(Hidden) Discard network LSAs.
- nssa**—(Hidden) Discard not-so-stubby area (NSSA) LSAs.
- opaque-area**—(Hidden) Discard opaque area-scope LSAs.
- purge**—(Hidden) Discard all entries in the link-state advertisement database. All link-state advertisements are set to **MAXAGE** and are flooded. The database is repopulated

when the originators of the link-state advertisements receive the **MAXAGE** link-state advertisements and reissue them.

realm (**ipv4-multicast** | **ipv4-unicast** | **ipv6-multicast**)—(OSPFv3 only) (Optional) Delete the entries for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

router—(Hidden) Discard router LSAs.

Required Privilege Level

clear

Related Documentation

- [show ospf database on page 711](#)
- [show ospf3 database on page 719](#)

List of Sample Output [clear ospf database on page 688](#)

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

Sample Output

clear ospf database

```
user@host> clear ospf database
```

clear (ospf | ospf3) io-statistics

Syntax	clear (ospf ospf3) io-statistics <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch and QFX Series)	clear (ospf ospf3) io-statistics
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Clear Open Shortest Path First (OSPF) input and output statistics.
Options	none —Clear OSPF input and output statistics. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	clear
List of Sample Output	clear ospf io-statistics on page 689
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ospf io-statistics

```
user@host> clear ospf io-statistics
```

clear (ospf | ospf3) neighbor

Syntax	clear (ospf ospf3) neighbor <area <i>area-id</i> > <instance <i>instance-name</i> > <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)> <neighbor> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)>
Syntax (EX Series Switch and QFX Series)	clear (ospf ospf3) neighbor <area <i>area-id</i> > <instance <i>instance-name</i> > <interface <i>interface-name</i> > <neighbor>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. realm option introduced in Junos OS Release 9.2. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Tear down Open Shortest Path First (OSPF) neighbor connections.
Options	none —Tear down OSPF connections with all neighbors for all routing instances. area <i>area-id</i> —(Optional) Tear down neighbor connections for the specified area only. instance <i>instance-name</i> —(Optional) Tear down neighbor connections for the specified routing instance only. interface <i>interface-name</i> —(Optional) Tear down neighbor connections for the specified interface only. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. neighbor —(Optional) Clear the state of the specified neighbor only. realm (ipv4-multicast ipv4-unicast ipv6-multicast) —(Optional) (OSPFv3 only) Clear the state of the specified OSPFv3 realm, or address family. Use the realm option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show (ospf ospf3) neighbor on page 741
List of Sample Output	clear ospf neighbor on page 691
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ospf neighbor

```
user@host> clear ospf neighbor
```

clear (ospf | ospf3) overload

Syntax	clear (ospf ospf3) overload <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	clear (ospf ospf3) overload <instance <i>instance-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Clear the Open Shortest Path First (OSPF) overload bit and rebuild link-state advertisements (LSAs).
Options	none —Clear the overload bit and rebuild LSAs for all routing instances. instance <i>instance-name</i> —(Optional) Clear the overload bit and rebuild LSAs for the specified routing instance only. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	clear
List of Sample Output	clear ospf overload on page 692
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ospf overload

```
user@host> clear ospf overload
```

clear (ospf | ospf3) statistics

Syntax	clear (ospf ospf3) statistics <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)>
Syntax (EX Series Switch and QFX Series)	clear (ospf ospf3) statistics <instance <i>instance-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. realm option introduced in Junos OS Release 9.2. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Clear Open Shortest Path First (OSPF) statistics.
Options	<p>none—Clear OSPF statistics.</p> <p>instance <i>instance-name</i>—(Optional) Clear statistics for the specified routing instance only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>realm (ipv4-multicast ipv4-unicast ipv6-multicast)—(Optional) (OSPFv3 only) Clear statistics for the specified OSPFv3 realm, or address family. Use the realm option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show (ospf ospf3) statistics on page 758
List of Sample Output	clear ospf statistics on page 693
Output Fields	See show (ospf ospf3) statistics for an explanation of output fields.

Sample Output

clear ospf statistics

The following sample output displays OSPF statistics before and after the **clear ospf statistics** command is entered:

```
user@host> show ospf statistics
```

Packet type	Total		Last 5 seconds	
	Sent	Received	Sent	Received
Hello	3254	2268	3	1
DbD	41	46	0	0

LSReq	8	7	0	0
LSUpdate	212	154	0	0
LSAck	65	98	0	0

DBDs retransmitted	:	3, last 5 seconds	:	0
LSAs flooded	:	12, last 5 seconds	:	0
LSAs flooded high-prio	:	0, last 5 seconds	:	0
LSAs retransmitted	:	0, last 5 seconds	:	0
LSAs transmitted to nbr:	:	3, last 5 seconds	:	0
LSAs requested	:	5, last 5 seconds	:	0
LSAs acknowledged	:	19, last 5 seconds	:	0

Flood queue depth	:	0
Total rexmit entries	:	0
db summaries	:	0
lsreq entries	:	0

Receive errors:
626 subnet mismatches

user@host> clear ospf statistics

user@host> show ospf statistics

Packet type	Total		Last 5 seconds	
	Sent	Received	Sent	Received
Hello	3	1	3	1
DbD	0	0	0	0
LSReq	0	0	0	0
LSUpdate	0	0	0	0
LSAck	0	0	0	0

DBDs retransmitted	:	0, last 5 seconds	:	0
LSAs flooded	:	0, last 5 seconds	:	0
LSAs flooded high-prio	:	0, last 5 seconds	:	0
LSAs retransmitted	:	0, last 5 seconds	:	0
LSAs transmitted to nbr:	:	0, last 5 seconds	:	0
LSAs requested	:	0, last 5 seconds	:	0
LSAs acknowledged	:	0, last 5 seconds	:	0

Flood queue depth	:	0
Total rexmit entries	:	0
db summaries	:	0
lsreq entries	:	0

Receive errors:
None

show (ospf | ospf3) backup coverage

Syntax	<pre>show (ospf ospf3) backup coverage <instance <i>instance-name</i>> < logical-system (all <i>logical-system-name</i>)> <realm (ipv4-unicast ipv6-unicast)> <topology <i>topology-name</i>></pre>
Syntax (QFX Series)	<pre>show (ospf ospf3) backup coverage <instance <i>instance-name</i>> <topology <i>topology-name</i>></pre>
Release Information	<p>Command introduced in Junos OS Release 10.0.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Display information about the level of backup coverage available for all the nodes and prefixes in the network.
Options	<p>none—Display information about the level backup coverage for all OSPF routing instances in all logical systems.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Display information about the level of backup coverage for all logical systems or for a specific logical system.</p> <p>instance <i>instance-name</i>—(Optional) Display information about the level of backup coverage for a specific OSPF routing instance.</p> <p>realm (ipv4-unicast ipv6-unicast)—(Optional) (OSPFv3 only) Display information about the level of backup coverage for the specific OSPFv3 realm, or address family.</p> <p>topology (default <i>topology-name</i>)—(Optional) (OSPFv2 only) Display information about the level of backup coverage for the specific OSPF topology.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> show (ospf ospf3) backup lsp on page 698
List of Sample Output	<p>show ospf backup coverage on page 696</p> <p>show ospf3 backup coverage on page 696</p>
Output Fields	<p>Table 106 on page 695 lists the output fields for the show (ospf ospf3) backup coverage command. Output fields are listed in the approximate order in which they appear.</p>

Table 106: show (ospf | ospf3) backup coverage Output Fields

Field Name	Field Description
Node Coverage	Information about backup coverage for each OSPF node.
Area	Area number. Area 0.0.0.0 is the backbone.

Table 106: show (ospf | ospf3) backup coverage Output Fields (*continued*)

Field Name	Field Description
Covered Nodes	Number of nodes for which backup coverage is available.
Total Nodes	Total number of OSPF nodes.
Route Coverage	Information about backup coverage for each type of OSPF route.
Path Type	Type of OSPF path: Intra , Inter , Ext1 , Ext2 , and All .
Covered Routes	For each path type, the number of routes for which backup coverage is available.
Total Routes	For each path type, the total number of configured routes.
Percent Covered	For all nodes and for each path type, the percentage for which backup coverage is available.

Sample Output

show ospf backup coverage

```

user@host> show ospf backup coverage
Topology default coverage:

Node Coverage:

Area              Covered  Total  Percent
                  Nodes   Nodes  Covered
0.0.0.0           4       5     80.00%

Route Coverage:

Path Type  Covered  Total  Percent
           Routes Routes  Covered
Intra      8       14    57.14%
Inter      0       0    100.00%
Ext1       0       0    100.00%
Ext2       1       1    100.00%
All        9       15    60.00%

```

show ospf3 backup coverage

```

user @host > show ospf3 backup coverage
show ospf3 backup coverage
Node Coverage:

Area              Covered  Total  Percent
                  Nodes   Nodes  Covered
0.0.0.0           4       5     80.00%

Route Coverage:

Path Type  Covered  Total  Percent
           Routes Routes  Covered

```

Intra	4	6	66.67%
Inter	0	0	100.00%
Ext1	0	0	100.00%
Ext2	1	1	100.00%
All	5	7	71.43%

show (ospf | ospf3) backup lsp

Syntax `show (ospf | ospf3) backup lsp`
`<logical-system (all | logical-system-name)>`
`<realm (ipv4-unicast | ipv6-unicast)>`

Release Information Command introduced in Junos OS Release 10.0.

Description Display information about MPLS label-switched-paths (LSPs) designated as backup routes for OSPF routes.



NOTE: MPLS LSPs can be used as backup routes only for routes in the default OSPFv2 topology and not for any configured topology. Additionally, MPLS LSPs cannot be used as backup routes for nondefault instances either for OSPFv2 or OSPFv3.

Options `none`—Display information all MPLS LSPs designated as backup routes.

`logical-system (all | logical-system-name)`—(Optional) Display information about MPLS LSPs designated as backup routes for all logical systems or a specific logical system.

`realm (ipv4-unicast | ipv6-unicast)`—(Optional) (OSPFv3 only) Display information about MPLS LSPs designated as backup routes for a specific realm, or address family.

Required Privilege Level view

Related Documentation • [show \(ospf | ospf3\) backup coverage on page 695](#)

List of Sample Output [show ospf backup lsp on page 699](#)
[show ospf3 backup lsp on page 699](#)

Output Fields [Table 107 on page 698](#) lists the output fields for the `show (ospf | ospf3) backup lsp` command. Output fields are listed in the approximate order in which they appear.

Table 107: show (ospf | ospf3) backup lsp Output Fields

Field Name	Field Description
<i>MPLS LSP name</i>	Name of each MPLS LSP designated as a backup path.
Egress	IP address of the egress router for the LSP.

Table 107: show (ospf | ospf3) backup lsp Output Fields (*continued*)

Field Name	Field Description
Status	<p>State of the LSP:</p> <ul style="list-style-type: none"> • Up—The router can detect RSVP hello messages from the neighbor. • Down—The router has received one of the following indications: <ul style="list-style-type: none"> • Communication failure from the neighbor. • Communication from IGP that the neighbor is unavailable. • Change in the sequence numbers in the RSVP hello messages sent by the neighbor. • Deleted—The LSP is no longer available as a backup path.
Last change	Time elapsed since the neighbor state changed either from up or down or from down to up . The format is <i>hh:mm:ss</i> .
TE-metric	Configured traffic engineering metric.
Metric	Configured metric.

Sample Output

show ospf backup lsp

```

user@host> show ospf backup lsp
tobanff
  Egress: 10.255.71.239, Status: up, Last change: 00:00:23
  TE-metric: 0, Metric: 0

```

Sample Output

show ospf3 backup lsp

```

user@host> show ospf3 backup lsp
tobanff
  Egress: 10.255.71.239, Status: up, Last change: 00:00:45
  TE-metric: 0, Metric: 0

```

show (ospf | ospf3) backup neighbor

Syntax	show (ospf ospf3) backup neighbor <area <i>area-id</i> > <instance (default <i>instance-name</i>)> <logical-system (default ipv4-multicast <i>logical-system-name</i>)> <topology (default ipv4-multicast <i>topology-name</i>)>
Syntax (QFX Series)	show (ospf ospf3) backup neighbor <area <i>area-id</i> > <instance <i>instance-name</i> > <topology (default ipv4-multicast <i>topology-name</i>)>
Release Information	Command introduced in Junos OS Release 10.0. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display the neighbors through which direct next hops for the backup paths are available.
Options	<p>none—Display all neighbors that have direct next hops for backup paths.</p> <p>area <i>area-id</i>—(Optional) Display the area information.</p> <p>instance (default <i>instance-name</i>)—(Optional) Display information about the default routing instance or a particular routing instance.</p> <p>logical-system (default ipv4-multicast <i>logical-system-name</i>)—(Optional) Display information about the default logical system, IPv4 multicast logical system, or a particular logical system.</p> <p>topology (default ipv4-multicast <i>topology-name</i>)—(OSPFv2 only) (Optional) Display information about the default topology, IPv4 multicast topology, or a particular topology.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show (ospf ospf3) backup spf on page 702
List of Sample Output	show ospf backup neighbor on page 701
Output Fields	Table 108 on page 700 lists the output fields for the show (ospf ospf3) backup neighbor command. Output fields are listed in the approximate order in which they appear.

Table 108: show (ospf | ospf3) backup neighbor Output Fields

Field Name	Field Description	Level of Output
Neighbor to Self Metric	Metric from the backup neighbor to the OSPF node.	All levels
Self to Neighbor Metric	Metric from the OSPF node to the backup neighbor.	All levels

Table 108: show (ospf |ospf3) backup neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
Direct next-hop	Interface and address of the direct next hop.	All levels

Sample Output

show ospf backup neighbor

```
user@host> show ospf backup neighbor
Topology default backup neighbors:

Area 0.0.0.5 backup neighbors:

10.0.0.5
  Neighbor to Self Metric: 5
  Self to Neighbor Metric: 5
  Direct next-hop: ge-4/0/0.111 via 10.0.175.5

10.0.0.6
  Neighbor to Self Metric: 5
  Self to Neighbor Metric: 5
  Direct next-hop: ge-4/1/0.110 via 10.0.176.6
```

show (ospf | ospf3) backup spf

Syntax	<code>show (ospf ospf3) backup spf</code> <code><brief detail></code> <code><area <i>area-id</i>></code> <code><instance <i>instance-name</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code> <code><no-coverage></code> <code><node-id></code> <code><realm (ipv4-unicast ipv6-unicast)></code> <code><topology (default ipv4-multicast <i>topology-name</i>)></code>
Syntax (QFX Series)	<code>show (ospf ospf3) backup spf</code> <code><brief detail></code> <code><area <i>area-id</i>></code> <code><instance <i>instance-name</i>></code> <code><no-coverage></code> <code><node-id></code> <code><topology (default ipv4-multicast <i>topology-name</i>)></code>
Release Information	Command introduced in JUNOS Release 10.0. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display information about OSPF shortest-path-first calculations for backup paths.
Options	<p>none—Display information about OSPF shortest-path-first (SPF) calculations for all backup paths for all destination nodes.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>area <i>area-id</i>—(Optional) Display the area information.</p> <p>instance <i>instance-name</i>—(Optional) Display information about the routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Display information about all logical systems or a specific logical system.</p> <p>no-coverage—(Optional) Display information if there is no backup coverage.</p> <p>node-id—(Optional) Display information about the node specified.</p> <p>realm (ipv4-unicast ipv6-unicast)—(Optional) Display information about the ipv4 or ipv6 realm.</p> <p>topology (default ipv4-multicast <i>topology-name</i>)—(Optional) (OSPFv2 only) Display information about the default topology, IPv4 multicast topology, or a specific topology.</p>
Required Privilege Level	view
List of Sample Output	show ospf backup spf on page 703 show ospf backup spf detail on page 703 show ospf3 backup spf on page 706

Output Fields Table 109 on page 703 lists the output fields for the **show (ospf |ospf3) backup spf** command. Output fields are listed in the approximate order in which they appear.

Table 109: show (ospf |ospf3) backup spf Output Fields

Field Name	Field Description	Level of Output
Area <i>area-id</i> results	Area for which the results are displayed. Area 0.0.0.0 is the backbone area.	All levels
<i>address</i>	Address of the node for which the results are displayed.	All levels
Self to Destination Metric	Metric from the node to the destination.	All levels
Parent Node	Address of the parent node.	All levels
Primary next-hop	Address of the next hop.	All levels
Backup Neighbor	Address of the backup neighbor or LSP endpoint and the following information: <ul style="list-style-type: none"> Neighbor to Destination Metric Neighbor to Self Metric Self to Neighbor Metric Status (Eligible, Not Eligible, Not Evaluated) and the reason for the status. <p>NOTE: If the backup neighbor is an LSP endpoint, it is indicated as such after the neighbor address.</p>	All levels

Sample Output

show ospf backup spf

```

user@host> show ospf backup spf
Topology default results:

Area 0.0.0.0 results:

pro16-d-1o0.xxx.yyyy.net
Self to Destination Metric: 1
Parent Node: pro16-b-1o0.xxx.yyyy.net
Primary next-hop: at-1/0/1.0
Backup Neighbor: pro16-c-1o0.xxx.yyyy.net (LSP endpoint)
  Neighbor to Destination Metric: 4, Neighbor to Self Metric: 3
  Self to Neighbor Metric: 3
  Not eligible, Reason: Path loops
Backup Neighbor: pro16-d-1o0.xxx.yyyy.net
  Neighbor to Destination Metric: 0, Neighbor to Self Metric: 1
  Self to Neighbor Metric: 1
  Not eligible, Reason: Primary next-hop link fate sharing
...

```

show ospf backup spf detail

```

user@host> show ospf backup spf detail

```

Topology default results:

Area 0.0.0.0 results:

11.14.10.2

Self to Destination Metric: 1
Parent Node: 10.255.70.103
Primary next-hop: ae0.0
Backup Neighbor: 10.255.71.243
Neighbor to Destination Metric: 2, Neighbor to Self Metric: 1
Self to Neighbor Metric: 1
Not eligible, Reason: Path loops
Backup Neighbor: 10.255.71.52
Neighbor to Destination Metric: 15, Neighbor to Self Metric: 15
Self to Neighbor Metric: 1
Not eligible, Reason: Primary next-hop link fate sharing
Backup Neighbor: 10.255.71.242
Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
Self to Neighbor Metric: 1
Not eligible, Reason: Path loops

10.255.71.52

Self to Destination Metric: 1
Parent Node: 11.14.10.2
Primary next-hop: ae0.0 via 11.14.10.2
Backup Neighbor: 10.255.71.52
Neighbor to Destination Metric: 0, Neighbor to Self Metric: 15
Self to Neighbor Metric: 1
Not eligible, Reason: Primary next-hop link fate sharing
Backup Neighbor: 10.255.71.243
Neighbor to Destination Metric: 2, Neighbor to Self Metric: 1
Self to Neighbor Metric: 1
Not eligible, Reason: Path loops
Backup Neighbor: 10.255.71.242
Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
Self to Neighbor Metric: 1
Not eligible, Reason: Path loops

10.255.71.242

Self to Destination Metric: 1
Parent Node: 10.255.70.103
Primary next-hop: as0.0
Backup Neighbor: 10.255.71.242
Neighbor to Destination Metric: 0, Neighbor to Self Metric: 15
Self to Neighbor Metric: 1
Not eligible, Reason: Primary next-hop link fate sharing
Backup Neighbor: 10.255.71.243
Neighbor to Destination Metric: 2, Neighbor to Self Metric: 1
Self to Neighbor Metric: 1
Not eligible, Reason: Path loops
Backup Neighbor: 10.255.71.52
Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
Self to Neighbor Metric: 1
Not eligible, Reason: Path loops

10.255.71.243

Self to Destination Metric: 1
Parent Node: 10.255.70.103
Primary next-hop: so-6/0/0.0
Backup Neighbor: 10.255.71.243
Neighbor to Destination Metric: 0, Neighbor to Self Metric: 1

```

    Self to Neighbor Metric: 1
    Not eligible, Reason: Primary next-hop link fate sharing
Backup Neighbor: 10.255.71.52
    Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not eligible, Reason: Path loops
Backup Neighbor: 10.255.71.242
    Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not eligible, Reason: Path loops

12.15.0.1
    Self to Destination Metric: 2
    Parent Node: 10.255.71.243
    Primary next-hop: so-6/0/0.0
    Backup next-hop: ae0.0 via 11.14.10.2
Backup Neighbor: 10.255.71.243
    Neighbor to Destination Metric: 1, Neighbor to Self Metric: 1
    Self to Neighbor Metric: 1
    Not eligible, Reason: Primary next-hop link fate sharing
Backup Neighbor: 10.255.71.52
    Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Eligible, Reason: Contributes backup next-hop
Backup Neighbor: 10.255.71.242
    Neighbor to Destination Metric: 17, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not evaluated, Reason: Interface is already covered

10.255.71.238
    Self to Destination Metric: 2
    Parent Node: 10.255.71.243
    Primary next-hop: so-6/0/0.0
    Backup next-hop: as0.0
Backup Neighbor: 10.255.71.243
    Neighbor to Destination Metric: 1, Neighbor to Self Metric: 1
    Self to Neighbor Metric: 1
    Not eligible, Reason: Primary next-hop link fate sharing
Backup Neighbor: 10.255.71.242
    Neighbor to Destination Metric: 15, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Eligible, Reason: Contributes backup next-hop
Backup Neighbor: 10.255.71.52
    Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not evaluated, Reason: Interface is already covered

10.255.71.239
    Self to Destination Metric: 2
    Parent Node: 12.15.0.1
    Primary next-hop: so-6/0/0.0
    Backup next-hop: ae0.0 via 11.14.10.2
Backup Neighbor: 10.255.71.243
    Neighbor to Destination Metric: 1, Neighbor to Self Metric: 1
    Self to Neighbor Metric: 1
    Not eligible, Reason: Primary next-hop link fate sharing
Backup Neighbor: 10.255.71.52
    Neighbor to Destination Metric: 15, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Eligible, Reason: Contributes backup next-hop
Backup Neighbor: 10.255.71.242

```

```
Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15  
Self to Neighbor Metric: 1  
Not evaluated, Reason: Interface is already covered
```

14.15.0.2

```
Self to Destination Metric: 3  
Parent Node: 10.255.71.239  
Primary next-hop: so-6/0/0.0  
Backup next-hop: ae0.0 via 11.14.10.2  
Backup Neighbor: 10.255.71.243  
Neighbor to Destination Metric: 2, Neighbor to Self Metric: 1  
Self to Neighbor Metric: 1  
Not eligible, Reason: Primary next-hop link fate sharing  
Backup Neighbor: 10.255.71.52  
Neighbor to Destination Metric: 15, Neighbor to Self Metric: 15  
Self to Neighbor Metric: 1  
Eligible, Reason: Contributes backup next-hop  
Backup Neighbor: 10.255.71.242  
Neighbor to Destination Metric: 17, Neighbor to Self Metric: 15  
Self to Neighbor Metric: 1  
Not evaluated, Reason: Interface is already covered
```

show ospf3 backup spf

```
user@host> show ospf3 backup spf  
Area 0.0.0.0 results:
```

10.255.71.52;0.0.0.5

```
Self to Destination Metric: 1  
Parent Node: 10.255.70.103  
Primary next-hop: ae0.0  
Backup Neighbor: 10.255.71.243  
Neighbor to Destination Metric: 2, Neighbor to Self Metric: 1  
Self to Neighbor Metric: 1  
Not eligible, Reason: Path loops  
Backup Neighbor: 10.255.71.52  
Neighbor to Destination Metric: 15, Neighbor to Self Metric: 15  
Self to Neighbor Metric: 1  
Not eligible, Reason: Primary next-hop link fate sharing  
Backup Neighbor: 10.255.71.242  
Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15  
Self to Neighbor Metric: 1  
Not eligible, Reason: Path loops
```

10.255.71.52

```
Self to Destination Metric: 1  
Parent Node: 10.255.71.52;0.0.0.5  
Primary next-hop: ae0.0 via fe80::290:69ff:fe0f:67f0  
Backup Neighbor: 10.255.71.52  
Neighbor to Destination Metric: 0, Neighbor to Self Metric: 15  
Self to Neighbor Metric: 1  
Not eligible, Reason: Primary next-hop link fate sharing  
Backup Neighbor: 10.255.71.243  
Neighbor to Destination Metric: 2, Neighbor to Self Metric: 1  
Self to Neighbor Metric: 1  
Not eligible, Reason: Path loops  
Backup Neighbor: 10.255.71.242  
Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15  
Self to Neighbor Metric: 1  
Not eligible, Reason: Path loops
```

```

10.255.71.242
  Self to Destination Metric: 1
  Parent Node: 10.255.70.103
  Primary next-hop: as0.0
  Backup Neighbor: 10.255.71.242
    Neighbor to Destination Metric: 0, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not eligible, Reason: Primary next-hop link fate sharing
  Backup Neighbor: 10.255.71.243
    Neighbor to Destination Metric: 2, Neighbor to Self Metric: 1
    Self to Neighbor Metric: 1
    Not eligible, Reason: Path loops
  Backup Neighbor: 10.255.71.52
    Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not eligible, Reason: Path loops

10.255.71.243
  Self to Destination Metric: 1
  Parent Node: 10.255.70.103
  Primary next-hop: so-6/0/0.0
  Backup Neighbor: 10.255.71.243
    Neighbor to Destination Metric: 0, Neighbor to Self Metric: 1
    Self to Neighbor Metric: 1
    Not eligible, Reason: Primary next-hop link fate sharing
  Backup Neighbor: 10.255.71.52
    Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not eligible, Reason: Path loops
  Backup Neighbor: 10.255.71.242
    Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not eligible, Reason: Path loops

10.255.71.243;0.0.0.2
  Self to Destination Metric: 2
  Parent Node: 10.255.71.243
  Primary next-hop: so-6/0/0.0
  Backup next-hop: ae0.0 via fe80::290:69ff:fe0f:67f0
  Backup Neighbor: 10.255.71.243
    Neighbor to Destination Metric: 1, Neighbor to Self Metric: 1
    Self to Neighbor Metric: 1
    Not eligible, Reason: Primary next-hop link fate sharing
  Backup Neighbor: 10.255.71.52
    Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Eligible, Reason: Contributes backup next-hop
  Backup Neighbor: 10.255.71.242
    Neighbor to Destination Metric: 17, Neighbor to Self Metric: 15
    Self to Neighbor Metric: 1
    Not evaluated, Reason: Interface is already covered

10.255.71.238
  Self to Destination Metric: 2
  Parent Node: 10.255.71.243
  Primary next-hop: so-6/0/0.0
  Backup next-hop: as0.0
  Backup Neighbor: 10.255.71.243
    Neighbor to Destination Metric: 1, Neighbor to Self Metric: 1
    Self to Neighbor Metric: 1
    Not eligible, Reason: Primary next-hop link fate sharing

```

```
Backup Neighbor: 10.255.71.242
  Neighbor to Destination Metric: 15, Neighbor to Self Metric: 15
  Self to Neighbor Metric: 1
  Eligible, Reason: Contributes backup next-hop
Backup Neighbor: 10.255.71.52
  Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
  Self to Neighbor Metric: 1
  Not evaluated, Reason: Interface is already covered

10.255.71.239
  Self to Destination Metric: 2
  Parent Node: 10.255.71.243;0.0.0.2
  Primary next-hop: so-6/0/0.0
  Backup next-hop: ae0.0 via fe80::290:69ff:fe0f:67f0
Backup Neighbor: 10.255.71.243
  Neighbor to Destination Metric: 1, Neighbor to Self Metric: 1
  Self to Neighbor Metric: 1
  Not eligible, Reason: Primary next-hop link fate sharing
Backup Neighbor: 10.255.71.52
  Neighbor to Destination Metric: 15, Neighbor to Self Metric: 15
  Self to Neighbor Metric: 1
  Eligible, Reason: Contributes backup next-hop
Backup Neighbor: 10.255.71.242
  Neighbor to Destination Metric: 16, Neighbor to Self Metric: 15
  Self to Neighbor Metric: 1
  Not evaluated, Reason: Interface is already covered

10.255.71.239;0.0.0.4
  Self to Destination Metric: 3
  Parent Node: 10.255.71.239
  Primary next-hop: so-6/0/0.0
  Backup next-hop: ae0.0 via fe80::290:69ff:fe0f:67f0
Backup Neighbor: 10.255.71.243
  Neighbor to Destination Metric: 2, Neighbor to Self Metric: 1
  Self to Neighbor Metric: 1
  Not eligible, Reason: Primary next-hop link fate sharing
Backup Neighbor: 10.255.71.52
  Neighbor to Destination Metric: 15, Neighbor to Self Metric: 15
  Self to Neighbor Metric: 1
  Eligible, Reason: Contributes backup next-hop
Backup Neighbor: 10.255.71.242
  Neighbor to Destination Metric: 17, Neighbor to Self Metric: 15
  Self to Neighbor Metric: 1
  Not evaluated, Reason: Interface is already covered
```

show ospf context-identifier

Syntax	<pre>show ospf context-identifier <brief detail> <area <i>area-id</i>> <context-id> <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show ospf context-identifier <brief detail> <area <i>area-id</i>> <context-id> <instance <i>instance-name</i>></pre>
Release Information	<p>Command introduced in Junos OS Release 10.4.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Display the context identifier information processed and advertised by Open Shortest Path First (OSPF) for egress protection.
Options	<p>none—Display information about all context identifiers.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>area <i>area-id</i>—(Optional) Display information about the context identifier for the specified area.</p> <p>context-id—(Optional) Display information about the specified context identifier.</p> <p>instance <i>instance-name</i>—(Optional) Display information about the context identifier for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> <i>egress-protection (Layer 2 circuit)</i> in the <i>Junos OS VPNs Library for Routing Devices</i> <i>egress-protection (MPLS)</i> in the <i>Junos OS VPNs Library for Routing Devices</i>
List of Sample Output	<p>show ospf context-identifier on page 710</p> <p>show ospf context-identifier detail on page 710</p>
Output Fields	<p>Table 110 on page 710 lists the output fields for the show ospf context-identifier command. Output fields are listed in the approximate order in which they appear.</p>

Table 110: show ospf context-identifier Output Fields

Field Name	Field Description	Level of Output
Context	IPv4 address that defines a protection pair. The context is manually configured on both primary and protector provider edge (PE) devices.	All levels
Status	State of the path: active or inactive .	All levels
Metric	Advertised OSPF metric.	All levels
Area	OSPF area number.	All levels
Other Advertisements	Other advertisements received by the OSPF node: <ul style="list-style-type: none"> • Advertising router—Address of the device that sent the advertisement. • Type—Type of OSPF path: inter-area and stub. • Metric—Advertised OSPF metric. • None—No additional advertisements were received by the OSPF node. 	detail

Sample Output

show ospf context-identifier

```
user@host> show ospf context-identifier
Context-id: 2.2.4.3
Status: active, Metric: 65534, PE role: protector, Area: 0.0.0.0
```

show ospf context-identifier detail

```
user@host> show ospf context-identifier detail
Context-id: 88.24.13.1
Status: inactive, Metric: 0, PE role: protector, Area: 0.0.0.13
Other Advertisements:
Advertising router: 8.8.8.103
Type: stub link
Metric: 65534
```


show ospf database

Syntax	<pre>show ospf database <brief detail extensive summary> <advertising-router (address self)> <area area-id> <asbrsummary> <external> <instance instance-name> <link-local> <logical-system (all logical-system-name)> <lsa-id lsa-id> <netsummary> <network> <nssa> <opaque-area> <router></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show ospf database <brief detail extensive summary> <advertising-router (address self)> <area area-id> <asbrsummary> <external> <instance instance-name> <link-local> <lsa-id lsa-id> <netsummary> <network> <nssa> <opaque-area> <router></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>advertising-router self (address self) option introduced in Junos OS Release 9.5.</p> <p>advertising-router self (address self) option introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Display the entries in the OSPF version 2 (OSPFv2) link-state database, which contains data about link-state advertisement (LSA) packets.
Options	<p>none—Display standard information about entries in the OSPFv2 link-state database for all routing instances.</p> <p>brief detail extensive summary—(Optional) Display the specified level of output.</p> <p>advertising-router (address self)—(Optional) Display the LSAs advertised either by a particular routing device or by this routing device.</p> <p>area area-id—(Optional) Display the LSAs in a particular area.</p>

asbrsummary—(Optional) Display summary AS boundary router LSA entries.

external—(Optional) Display external LSAs.

instance *instance-name*—(Optional) Display all OSPF database information under the named routing instance.

link-local—(Optional) Display information about link-local LSAs.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

lsa-id *lsa-id*—(Optional) Display the LSA with the specified LSA identifier.

netsummary—(Optional) Display summary network LSAs.

network—(Optional) Display information about network LSAs.

nssa—(Optional) Display information about not-so-stubby area (NSSA) LSAs.

opaque-area—(Optional) Display opaque area-scope LSAs.

router—(Optional) Display information about router LSAs.

Required Privilege Level

view

Related Documentation

- [clear \(ospf | ospf3\) database on page 686](#)

List of Sample Output

[show ospf database on page 714](#)
[show ospf database brief on page 714](#)
[show ospf database detail on page 714](#)
[show ospf database extensive on page 716](#)
[show ospf database summary on page 718](#)

Output Fields

[Table 111 on page 712](#) describes the output fields for the **show ospf database** command. Output fields are listed in the approximate order in which they appear.

Table 111: show ospf database Output Fields

Field Name	Field Description	Level of Output
area	Area number. Area 0.0.0.0 is the backbone area.	All levels
Type	Type of link advertisement: ASBRSum , Extern , Network , NSSA , OpaqArea , Router , or Summary .	All levels
ID	LSA identifier included in the advertisement. An asterisk preceding the identifier marks database entries that originated from the local routing device.	All levels
Adv Rtr	Address of the routing device that sent the advertisement.	All levels
Seq	Link sequence number of the advertisement.	All levels

Table 111: show ospf database Output Fields (*continued*)

Field Name	Field Description	Level of Output
Age	Time elapsed since the LSA was originated, in seconds.	All levels
Opt	Optional OSPF capabilities associated with the LSA.	All levels
Cksum	Checksum value of the LSA.	All levels
Len	Length of the advertisement, in bytes.	All levels
Router	Router link-state advertisement information: <ul style="list-style-type: none"> bits—Flags describing the routing device that generated the LSP. link count—Number of links in the advertisement. id—ID of a routing device or subnet on the link. data—For stub networks, the subnet mask. Otherwise, the IP address of the routing device that generated the LSP. type—Type of link. It can be PointToPoint, Transit, Stub, or Virtual. TOS count—Number of type-of-service (ToS) entries in the advertisement. TOS 0 metric—Metric for ToS 0. TOS—Type-of-service (ToS) value. metric—Metric for the ToS. 	detail extensive
Network	Network link-state advertisement information: <ul style="list-style-type: none"> mask—Network mask. attached router—ID of the attached neighbor. 	detail extensive
Summary	Summary link-state advertisement information: <ul style="list-style-type: none"> mask—Network mask. TOS—Type-of-service (ToS) value. metric—Metric for the ToS. 	detail extensive
Gen timer	How long until the LSA is regenerated.	extensive
Aging timer	How long until the LSA expires.	extensive
Installed <i>hh:mm:ss</i> ago	How long ago the route was installed.	extensive
expires in <i>hh:mm:ss</i>	How long until the route expires.	extensive
sent <i>hh:mm:ss</i> ago	How long ago the LSA was sent.	extensive
Last changed <i>hh:mm:ss</i> ago	How long ago the route was changed.	extensive
Change count	Number of times the route has changed.	extensive

Table 111: show ospf database Output Fields (*continued*)

Field Name	Field Description	Level of Output
Ours	Indicates that this is a local advertisement.	extensive
Router LSAs	Number of router link-state advertisements in the link-state database.	summary
Network LSAs	Number of network link-state advertisements in the link-state database.	summary
Summary LSAs	Number of summary link-state advertisements in the link-state database.	summary
NSSA LSAs	Number of not-so-stubby area link-state advertisements in the link-state database.	summary

Sample Output

show ospf database

```

user@host> show ospf database
OSPF link state database, Area 0.0.0.1
  Type      ID            Adv Rtr      Seq          Age    Opt  Cksum  Len
Router     10.255.70.103  10.255.70.103 0x80000002   215    0x20 0x4112  48
Router     *10.255.71.242 10.255.71.242 0x80000002   214    0x20 0x11b1  48
Summary    *23.1.1.0      10.255.71.242 0x80000002   172    0x20 0x6d72  28
Summary    *24.1.1.0      10.255.71.242 0x80000002   177    0x20 0x607e  28
NSSA       *33.1.1.1      10.255.71.242 0x80000002   217    0x28 0x73bd  36

      OSPF link state database, Area 0.0.0.2
  Type      ID            Adv Rtr      Seq          Age    Opt  Cksum  Len
Router     10.255.71.52   10.255.71.52   0x80000004   174    0x20 0xd021  36
Router     *10.255.71.242 10.255.71.242 0x80000003   173    0x20 0xe191  36
Network    *23.1.1.1      10.255.71.242 0x80000002   173    0x20 0x9c76  32
Summary    *12.1.1.0      10.255.71.242 0x80000001   217    0x20 0xfeec  28
Summary    *24.1.1.0      10.255.71.242 0x80000002   177    0x20 0x607e  28
NSSA       *33.1.1.1      10.255.71.242 0x80000001   222    0x28 0xe047  36

      OSPF link state database, Area 0.0.0.3
  Type      ID            Adv Rtr      Seq          Age    Opt  Cksum  Len
Router     10.255.71.238   10.255.71.238 0x80000003   179    0x20 0x3942  36
Router     *10.255.71.242 10.255.71.242 0x80000003   177    0x20 0xf37d  36
Network    *24.1.1.1      10.255.71.242 0x80000002   177    0x20 0xc591  32
Summary    *12.1.1.0      10.255.71.242 0x80000001   217    0x20 0xfeec  28
Summary    *23.1.1.0      10.255.71.242 0x80000002   172    0x20 0x6d72  28
NSSA       *33.1.1.1      10.255.71.242 0x80000001   222    0x28 0xeb3b  36

```

show ospf database brief

The output for the **show ospf database brief** command is identical to that for the **show ospf database** command. For sample output, see [show ospf database on page 714](#).

show ospf database detail

```

user@host> show ospf database detail
OSPF link state database, Area 0.0.0.1
  Type      ID            Adv Rtr      Seq          Age    Opt  Cksum  Len
Router     10.255.70.103  10.255.70.103 0x80000002   261    0x20 0x4112  48

```

```

bits 0x0, link count 2
id 10.255.71.242, data 12.1.1.1, Type PointToPoint (1)
TOS count 0, TOS 0 metric 1
id 12.1.1.0, data 255.255.255.0, Type Stub (3)
TOS count 0, TOS 0 metric 1
Router *10.255.71.242 10.255.71.242 0x80000002 260 0x20 0x11b1 48
bits 0x3, link count 2
id 10.255.70.103, data 12.1.1.2, Type PointToPoint (1)
TOS count 0, TOS 0 metric 1
id 12.1.1.0, data 255.255.255.0, Type Stub (3)
TOS count 0, TOS 0 metric 1
Summary *23.1.1.0 10.255.71.242 0x80000002 218 0x20 0x6d72 28
mask 255.255.255.0
TOS 0x0, metric 1
Summary *24.1.1.0 10.255.71.242 0x80000002 223 0x20 0x607e 28
mask 255.255.255.0
TOS 0x0, metric 1
NSSA *33.1.1.1 10.255.71.242 0x80000002 263 0x28 0x73bd 36
mask 255.255.255.255
Type 2, TOS 0x0, metric 0, fwd addr 12.1.1.2, tag 0.0.0.0

```

OSPF link state database, Area 0.0.0.2

Type	ID	Adv Rtr	Seq	Age	Opt	Cksum	Len
Router	10.255.71.52	10.255.71.52	0x80000004	220	0x20	0xd021	36
bits 0x0, link count 1							
id 23.1.1.1, data 23.1.1.2, Type Transit (2)							
TOS count 0, TOS 0 metric 1							
Router	*10.255.71.242	10.255.71.242	0x80000003	219	0x20	0xe191	36
bits 0x3, link count 1							
id 23.1.1.1, data 23.1.1.1, Type Transit (2)							
TOS count 0, TOS 0 metric 1							
Network	*23.1.1.1	10.255.71.242	0x80000002	219	0x20	0x9c76	32
mask 255.255.255.0							
attached router 10.255.71.242							
attached router 10.255.71.52							
Summary	*12.1.1.0	10.255.71.242	0x80000001	263	0x20	0xfeec	28
mask 255.255.255.0							
TOS 0x0, metric 1							
Summary	*24.1.1.0	10.255.71.242	0x80000002	223	0x20	0x607e	28
mask 255.255.255.0							
TOS 0x0, metric 1							
NSSA	*33.1.1.1	10.255.71.242	0x80000001	268	0x28	0xe047	36
mask 255.255.255.255							
Type 2, TOS 0x0, metric 0, fwd addr 23.1.1.1, tag 0.0.0.0							

OSPF link state database, Area 0.0.0.3

Type	ID	Adv Rtr	Seq	Age	Opt	Cksum	Len
Router	10.255.71.238	10.255.71.238	0x80000003	225	0x20	0x3942	36
bits 0x0, link count 1							
id 24.1.1.1, data 24.1.1.2, Type Transit (2)							
TOS count 0, TOS 0 metric 1							
Router	*10.255.71.242	10.255.71.242	0x80000003	223	0x20	0xf37d	36
bits 0x3, link count 1							
id 24.1.1.1, data 24.1.1.1, Type Transit (2)							
TOS count 0, TOS 0 metric 1							
Network	*24.1.1.1	10.255.71.242	0x80000002	223	0x20	0xc591	32
mask 255.255.255.0							
attached router 10.255.71.242							
attached router 10.255.71.238							
Summary	*12.1.1.0	10.255.71.242	0x80000001	263	0x20	0xfeec	28
mask 255.255.255.0							

```

TOS 0x0, metric 1
Summary *23.1.1.0      10.255.71.242    0x80000002    218  0x20 0x6d72  28
mask 255.255.255.0
TOS 0x0, metric 1
NSSA  *33.1.1.1      10.255.71.242    0x80000001    268  0x28 0xeb3b  36
mask 255.255.255.255
Type 2, TOS 0x0, metric 0, fwd addr 24.1.1.1, tag 0.0.0.0

```

show ospf database extensive

```

user@host> show ospf database extensive
  OSPF link state database, Area 0.0.0.1
  Type      ID          Adv Rtr          Seq      Age  Opt  Cksum  Len
Router  10.255.70.103    10.255.70.103    0x80000002   286  0x20 0x4112  48
  bits 0x0, link count 2
  id 10.255.71.242, data 12.1.1.1, Type PointToPoint (1)
  TOS count 0, TOS 0 metric 1
  id 12.1.1.0, data 255.255.255.0, Type Stub (3)
  TOS count 0, TOS 0 metric 1
  Aging timer 00:55:14
  Installed 00:04:43 ago, expires in 00:55:14
  Last changed 00:04:43 ago, Change count: 2
Router *10.255.71.242  10.255.71.242    0x80000002   285  0x20 0x11b1  48
  bits 0x3, link count 2
  id 10.255.70.103, data 12.1.1.2, Type PointToPoint (1)
  TOS count 0, TOS 0 metric 1
  id 12.1.1.0, data 255.255.255.0, Type Stub (3)
  TOS count 0, TOS 0 metric 1
  Gen timer 00:45:15
  Aging timer 00:55:15
  Installed 00:04:45 ago, expires in 00:55:15, sent 00:04:43 ago
  Last changed 00:04:45 ago, Change count: 2, Ours
Summary *23.1.1.0      10.255.71.242    0x80000002   243  0x20 0x6d72  28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:57
Aging timer 00:55:57
Installed 00:04:03 ago, expires in 00:55:57, sent 00:04:01 ago
Last changed 00:04:48 ago, Change count: 1, Ours
Summary *24.1.1.0      10.255.71.242    0x80000002   248  0x20 0x607e  28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:52
Aging timer 00:55:52
Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:06 ago
Last changed 00:04:48 ago, Change count: 1, Ours
NSSA  *33.1.1.1      10.255.71.242    0x80000002   288  0x28 0x73bd  36
mask 255.255.255.255
Type 2, TOS 0x0, metric 0, fwd addr 12.1.1.2, tag 0.0.0.0
Gen timer 00:45:12
Aging timer 00:55:12
Installed 00:04:48 ago, expires in 00:55:12, sent 00:04:48 ago
Last changed 00:04:48 ago, Change count: 2, Ours

  OSPF link state database, Area 0.0.0.2
  Type      ID          Adv Rtr          Seq      Age  Opt  Cksum  Len
Router  10.255.71.52    10.255.71.52    0x80000004   245  0x20 0xd021  36
  bits 0x0, link count 1
  id 23.1.1.1, data 23.1.1.2, Type Transit (2)
  TOS count 0, TOS 0 metric 1
  Aging timer 00:55:55

```

```

    Installed 00:04:02 ago, expires in 00:55:55
    Last changed 00:04:02 ago, Change count: 2
Router *10.255.71.242    10.255.71.242    0x80000003    244    0x20 0xe191    36
    bits 0x3, link count 1
    id 23.1.1.1, data 23.1.1.1, Type Transit (2)
    TOS count 0, TOS 0 metric 1
    Gen timer 00:45:56
    Aging timer 00:55:56
    Installed 00:04:04 ago, expires in 00:55:56, sent 00:04:02 ago
    Last changed 00:04:04 ago, Change count: 2, Ours
Network *23.1.1.1        10.255.71.242    0x80000002    244    0x20 0x9c76    32
    mask 255.255.255.0
    attached router 10.255.71.242
    attached router 10.255.71.52
    Gen timer 00:45:56
    Aging timer 00:55:56
    Installed 00:04:04 ago, expires in 00:55:56, sent 00:04:02 ago
    Last changed 00:04:04 ago, Change count: 1, Ours
Summary *12.1.1.0        10.255.71.242    0x80000001    288    0x20 0xfeec    28
    mask 255.255.255.0
    TOS 0x0, metric 1
    Gen timer 00:45:12
    Aging timer 00:55:12
    Installed 00:04:48 ago, expires in 00:55:12, sent 00:04:04 ago
    Last changed 00:04:48 ago, Change count: 1, Ours
Summary *24.1.1.0        10.255.71.242    0x80000002    248    0x20 0x607e    28
    mask 255.255.255.0
    TOS 0x0, metric 1
    Gen timer 00:45:52
    Aging timer 00:55:52
    Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:04 ago
    Last changed 00:04:48 ago, Change count: 1, Ours
NSSA  *33.1.1.1        10.255.71.242    0x80000001    293    0x28 0xe047    36
    mask 255.255.255.255
    Type 2, TOS 0x0, metric 0, fwd addr 23.1.1.1, tag 0.0.0.0
    Gen timer 00:45:07
    Aging timer 00:55:07
    Installed 00:04:53 ago, expires in 00:55:07, sent 00:04:04 ago
    Last changed 00:04:53 ago, Change count: 1, Ours

```

```

    OSPF link state database, Area 0.0.0.3
Type      ID          Adv Rtr      Seq      Age  Opt  Cksum  Len
Router  10.255.71.238    10.255.71.238    0x80000003    250  0x20 0x3942    36
    bits 0x0, link count 1
    id 24.1.1.1, data 24.1.1.2, Type Transit (2)
    TOS count 0, TOS 0 metric 1
    Aging timer 00:55:50
    Installed 00:04:07 ago, expires in 00:55:50
    Last changed 00:04:07 ago, Change count: 2
Router *10.255.71.242    10.255.71.242    0x80000003    248    0x20 0xf37d    36
    bits 0x3, link count 1
    id 24.1.1.1, data 24.1.1.1, Type Transit (2)
    TOS count 0, TOS 0 metric 1
    Gen timer 00:45:52
    Aging timer 00:55:52
    Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:06 ago
    Last changed 00:04:08 ago, Change count: 2, Ours
Network *24.1.1.1        10.255.71.242    0x80000002    248    0x20 0xc591    32
    mask 255.255.255.0
    attached router 10.255.71.242
    attached router 10.255.71.238

```

```
Gen timer 00:45:52
Aging timer 00:55:52
Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:06 ago
Last changed 00:04:08 ago, Change count: 1, Ours
Summary *12.1.1.0      10.255.71.242    0x80000001    288  0x20 0xfeec  28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:12
Aging timer 00:55:12
Installed 00:04:48 ago, expires in 00:55:12, sent 00:04:13 ago
Last changed 00:04:48 ago, Change count: 1, Ours
Summary *23.1.1.0      10.255.71.242    0x80000002    243  0x20 0x6d72  28
mask 255.255.255.0
TOS 0x0, metric 1
Gen timer 00:45:57
Aging timer 00:55:57
Installed 00:04:03 ago, expires in 00:55:57, sent 00:04:01 ago
Last changed 00:04:48 ago, Change count: 1, Ours
NSSA  *33.1.1.1        10.255.71.242    0x80000001    293  0x28 0xeb3b  36
mask 255.255.255.255
Type 2, TOS 0x0, metric 0, fwd addr 24.1.1.1, tag 0.0.0.0
Gen timer 00:45:07
Aging timer 00:55:07
Installed 00:04:53 ago, expires in 00:55:07, sent 00:04:13 ago
Last changed 00:04:53 ago, Change count: 1, Ours
```

show ospf database summary

```
user@host> show ospf database summary
Area 0.0.0.1:
  2 Router LSAs
  2 Summary LSAs
  1 NSSA LSAs
Area 0.0.0.2:
  2 Router LSAs
  1 Network LSAs
  2 Summary LSAs
  1 NSSA LSAs
Area 0.0.0.3:
  2 Router LSAs
  1 Network LSAs
  2 Summary LSAs
  1 NSSA LSAs
Externals:
Interface fe-2/2/1.0:
Interface ge-0/3/2.0:
Interface so-0/1/2.0:
Interface so-0/1/2.0:
```


show ospf3 database

Syntax	<pre>show ospf3 database <brief detail extensive summary> <advertising-router (address self)> <area area-id> <external> <instance instance-name> <inter-area-prefix> <inter-area-router> <intra-area-prefix> <link> <link-local> <logical-system (all logical-system-name)> <lsa-id lsa-id> <network> <nssa> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)> <router></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show ospf3 database <brief detail extensive summary> <advertising-router (address self)> <area area-id> <external> <instance instance-name> <inter-area-prefix> <inter-area-router> <intra-area-prefix> <link> <link-local> <lsa-id lsa-id> <network> <nssa> <router></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>realm option introduced in Junos OS Release 9.2.</p> <p>advertising-router (address self) option introduced in Junos Release 9.5.</p> <p>advertising-router (address self) option introduced in Junos OS Release 9.5 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Display the entries in the OSPF version 3 (OSPFv3) link-state database, which contains data about link-state advertisement (LSA) packets.
Options	<p>none—Display standard information about all entries in the OSPFv3 link-state database.</p> <p>brief detail extensive summary—(Optional) Display the specified level of output.</p> <p>advertising-router (address self)—(Optional) Display the LSAs advertised either by a particular routing device or by this routing device.</p>

area *area-id*—(Optional) Display the LSAs in a particular area.

external—(Optional) Display external LSAs.

instance *instance-name*—(Optional) Display all OSPF database information under the named routing instance.

inter-area-prefix—(Optional) Display information about interarea-prefix LSAs.

inter-area-router—(Optional) Display information about interarea-router LSAs.

intra-area-prefix—(Optional) Display information about intra-area-prefix LSAs.

link—(Optional) Display information about link LSAs.

link-local—(Optional) Display information about link-local LSAs.

logical-system (**all** | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

lsa-id *lsa-id*—(Optional) Display the LSA with the specified LSA identifier.

network—(Optional) Display information about network LSAs.

nssa—(Optional) Display information about not-so-stubby area (NSSA) LSAs.

realm (**ipv4-multicast** | **ipv4-unicast** | **ipv6-multicast**)—(Optional) Display information about the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family other than IPv6 unicast, which is the default.

router—(Optional) Display information about router LSAs.

Required Privilege Level view

Related Documentation • [clear \(ospf | ospf3\) database on page 686](#)

List of Sample Output [show ospf3 database brief on page 725](#)
[show ospf3 database extensive on page 725](#)
[show ospf3 database summary on page 728](#)

Output Fields [Table 112 on page 720](#) lists the output fields for the **show ospf3 database** command. Output fields are listed in the approximate order in which they appear.

Table 112: show ospf3 database Output Fields

Field Name	Field Description	Level of Output
OSPF link state database, area <i>area-number</i>	Entries in the link-state database for this area.	brief detail extensive
OSPF AS SCOPE link state database	Entries in the AS scope link-state database.	brief detail extensive

Table 112: show ospf3 database Output Fields (*continued*)

Field Name	Field Description	Level of Output
OSPF Link-Local link state database, interface <i>interface-name</i>	Entries in the link-local link-state database for this interface.	brief detail extensive
area	Area number. Area 0.0.0.0 is the backbone area.	All levels
Type	Type of link advertisement: Extern , InterArPfx , InterArRtr , IntraArPrx , Link , Network , NSSA , or Router .	brief detail extensive
ID	Link identifier included in the advertisement. An asterisk (*) preceding the identifier marks database entries that originated from the local routing device.	brief detail extensive
Adv Rtr	Address of the routing device that sent the advertisement.	brief detail extensive
Seq	Link sequence number of the advertisement.	brief detail extensive
Age	Time elapsed since the LSA was originated, in seconds.	brief detail extensive
Cksum	Checksum value of the LSA.	brief detail extensive
Len	Length of the advertisement, in bytes.	brief detail extensive
Router (Router Link-State Advertisements)		
bits	Flags describing the routing device that generated the LSP.	detail extensive
Options	Option bits carried in the router LSA.	detail extensive
For Each Router Link		
Type	Type of interface. The value of all other output fields describing a routing device interface depends on the interface's type: <ul style="list-style-type: none"> • PointToPoint (1)—Point-to-point connection to another routing device. • Transit (2)—Connection to a transit network. • Virtual (4)—Virtual link. 	detail extensive
Loc-if-id	Local interface ID assigned to the interface that uniquely identifies the interface with the routing device.	detail extensive
Nbr-if-id	Interface ID of the neighbor's interface for this routing device link.	detail extensive
Nbr-rtr-id	Router ID of the neighbor routing device (for type 2 interfaces, the attached link's designated router).	detail extensive
Metric	Cost of the router link.	detail extensive
Gen timer	How long until the LSA is regenerated, in the format <i>hours:minutes:seconds</i> .	extensive

Table 112: show ospf3 database Output Fields (*continued*)

Field Name	Field Description	Level of Output
Aging timer	How long until the LSA expires, in the format <i>hours:minutes:seconds</i> .	extensive
Installed <i>nn:nn:nn</i> ago	How long ago the route was installed, in the format <i>hours:minutes:seconds</i> .	extensive
expires in <i>nn:nn:nn</i>	How long until the route expires, in the format <i>hours:minutes:seconds</i> .	extensive
sent <i>nn:nn:nn</i> ago	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format <i>hours:minutes:seconds</i> .	extensive
Ours	Indicates that this is a local advertisement.	extensive
Network (Network Link-State Advertisements)		
Options	Option bits carried in the network LSA.	detail extensive
Attached Router	Router IDs of each of the routing devices attached to the link. Only routing devices that are fully adjacent to the designated router are listed. The designated router includes itself in this list.	detail extensive
InterArPfx (Interarea-Prefix Link-State Advertisements)		
Prefix	IPv6 address prefix.	detail extensive
Prefix-options	Option bit associated with the prefix.	detail extensive
Metric	Cost of this route. Expressed in the same units as the interface costs in the router LSAs. When the interarea-prefix LSA is describing a route to a range of addresses, the cost is set to the maximum cost to any reachable component of the address range.	detail extensive
Gen timer	How long until the LSA is regenerated, in the format <i>hours:minutes:seconds</i> .	extensive
Aging timer	How long until the LSA expires, in the format <i>hours:minutes:seconds</i> .	extensive
Installed <i>nn:nn:nn</i> ago	How long ago the route was installed, in the format <i>hours:minutes:seconds</i> .	extensive
expires in <i>nn:nn:nn</i>	How long until the route expires, in the format <i>hours:minutes:seconds</i> .	extensive
sent <i>nn:nn:nn</i> ago	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format <i>hours:minutes:seconds</i> .	extensive
Ours	Indicates that this is a local advertisement.	extensive
InterArRtr (Interarea-Router Link-State Advertisements)		
Dest-router-id	Router ID of the routing device described by the LSA.	detail extensive
options	Optional capabilities supported by the routing device.	detail extensive

Table 112: show ospf3 database Output Fields (*continued*)

Field Name	Field Description	Level of Output
Metric	Cost of this route. Expressed in the same units as the interface costs in the router LSAs. When the interarea-prefix LSA is describing a route to a range of addresses, the cost is set to the maximum cost to any reachable component of the address range.	detail extensive
Prefix	IPv6 address prefix.	extensive
Prefix-options	Option bit associated with the prefix.	extensive
Extern (External Link-State Advertisements)		
Prefix	IPv6 address prefix.	detail extensive
Prefix-options	Option bit associated with the prefix.	detail extensive
Metric	Cost of the route, which depends on the value of Type .	detail extensive
Type <i>n</i>	Type of external metric: Type 1 or Type 2 .	detail extensive
Aging timer	How long until the LSA expires, in the format <i>hours:minutes:seconds</i> .	extensive
Installed <i>nn:nn:nn</i> ago	How long ago the route was installed, in the format <i>hours:minutes:seconds</i> .	extensive
expires in <i>nn:nn:nn</i>	How long until the route expires, in the format <i>hours:minutes:seconds</i> .	extensive
sent <i>nn:nn:nn</i> ago	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format <i>hours:minutes:seconds</i> .	extensive
Link (Link-State Advertisements)		
IPv6-Address	IPv6 link-local address on the link for which this link LSA originated.	detail extensive
Options	Option bits carried in the link LSA.	detail extensive
priority	Router priority of the interface attaching the originating routing device to the link.	detail extensive
Prefix-count	Number of IPv6 address prefixes contained in the LSA. The rest of the link LSA contains a list of IPv6 prefixes to be associated with the link.	detail extensive
Prefix	IPv6 address prefix.	detail extensive
Prefix-options	Option bit associated with the prefix.	detail extensive
Gen timer	How long until the LSA is regenerated, in the format <i>hours:minutes:seconds</i> .	extensive
Aging timer	How long until the LSA expires, in the format <i>hours:minutes:seconds</i> .	extensive

Table 112: show ospf3 database Output Fields (*continued*)

Field Name	Field Description	Level of Output
Installed <i>nn:nn:nn</i> ago	How long ago the route was installed, in the format <i>hours:minutes:seconds</i> .	extensive
expires in <i>nn:nn:nn</i>	How long until the route expires, in the format <i>hours:minutes:seconds</i> .	extensive
sent <i>nn:nn:nn</i> ago	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format <i>hours:minutes:seconds</i> .	extensive
Ours	Indicates that this is a local advertisement.	extensive
IntraArPfx (Intra-Area-Prefix Link-State Advertisements)		
Ref-lsa-type	LSA type of the referenced LSA. <ul style="list-style-type: none"> Router—Address prefixes are associated with a router LSA. Network—Address prefixes are associated with a network LSA. 	detail extensive
Ref-lsa-id	Link-state ID of the referenced LSA.	detail extensive
Ref-router-id	Advertising router ID of the referenced LSA.	detail extensive
Prefix-count	Number of IPv6 address prefixes contained in the LSA. The rest of the link LSA contains a list of IPv6 prefixes to be associated with the link.	detail extensive
Prefix	IPv6 address prefix.	detail extensive
Prefix-options	Option bit associated with the prefix.	detail extensive
Metric	Cost of this prefix. Expressed in the same units as the interface costs in the router LSAs.	detail extensive
Gen timer	How long until the LSA is regenerated, in the format <i>hours:minutes:seconds</i> .	extensive
Aging timer	How long until the LSA expires, in the format <i>hours:minutes:seconds</i> .	extensive
Installed <i>hh:mm:ss</i> ago	How long ago the route was installed, in the format <i>hours:minutes:seconds</i> .	extensive
expires in <i>hh:mm:ss</i>	How long until the route expires, in the format <i>hours:minutes:seconds</i> .	extensive
sent <i>hh:mm:ss</i> ago	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format <i>hours:minutes:seconds</i> .	extensive
<i>n</i> Router LSAs	Number of router LSAs in the link-state database.	summary
<i>n</i> Network LSAs	Number of network LSAs in the link-state database.	summary
<i>n</i> InterArPfx LSAs	Number of interarea-prefix LSAs in the link-state database.	summary

Table 112: show ospf3 database Output Fields (*continued*)

Field Name	Field Description	Level of Output
<i>n</i> InterArRtr LSAs	Number of interarea-router LSAs in the link-state database.	summary
<i>n</i> IntraArPfx LSAs	Number of intra-area-prefix LSAs in the link-state database.	summary
Externals	Display of the external LSA database.	summary
<i>n</i> Extern LSAs	Number of external LSAs in the link-state database.	summary
Interface <i>interface-name</i>	Name of the interface for which link-local LSA information is displayed.	summary
<i>n</i> Link LSAs	Number of link LSAs in the link-state database.	summary

Sample Output

show ospf3 database brief

```

user@host> show ospf3 database brief
      OSPF3 link state database, area 0.0.0.0
      Type      ID          Adv Rtr      Seq          Age    Cksum  Len
      Router    0.0.0.1      10.255.4.85  0x80000003   885    0xa697  40
      Router    *0.0.0.1     10.255.4.93  0x80000002   953    0xc677  40
      InterArPfx *0.0.0.2     10.255.4.93  0x80000001   910    0xb96f  44
      InterArRtr *0.0.0.1     10.255.4.93  0x80000001   910    0xe159  32
      IntraArPfx *0.0.0.1     10.255.4.93  0x80000002   432    0x788f  72

      OSPF3 link state database, area 0.0.0.1
      Type      ID          Adv Rtr      Seq          Age    Cksum  Len
      Router    *0.0.0.1     10.255.4.93  0x80000003   916    0xea40  40
      Router    0.0.0.1     10.255.4.97  0x80000006   851    0xc95b  40
      Network    0.0.0.2     10.255.4.97  0x80000002   916    0x4598  32
      InterArPfx *0.0.0.1     10.255.4.93  0x80000002   117    0xa980  44
      InterArPfx *0.0.0.2     10.255.4.93  0x80000002    62    0xd47e  44
      NSSA       0.0.0.1     10.255.4.97  0x80000002   362    0x45ee  44
      IntraArPfx 0.0.0.1     10.255.4.97  0x80000006   851    0x2f77  52

      OSPF3 AS SCOPE link state database
      Type      ID          Adv Rtr      Seq          Age    Cksum  Len
      Extern     0.0.0.1     10.255.4.85  0x80000002    63    0x9b86  44
      Extern     *0.0.0.1     10.255.4.93  0x80000001   910    0x59c9  44

      OSPF3 Link-Local link state database, interface ge-1/3/0.0
      Type      ID          Adv Rtr      Seq          Age    Cksum  Len
      Link       *0.0.0.2     10.255.4.93  0x80000003   916    0x4dab  64

```

show ospf3 database extensive

```

user@host> show ospf3 database extensive
      OSPF3 link state database, area 0.0.0.0
      Type      ID          Adv Rtr      Seq          Age    Cksum  Len
      Router    0.0.0.1     10.255.4.85  0x80000003  1028    0xa697  40
      bits 0x2, Options 0x13
      Type PointToPoint (1), Metric 10

```

```

    Loc-If-Id 2, Nbr-If-Id 3, Nbr-Rtr-Id 10.255.4.93
    Aging timer 00:42:51
    Installed 00:17:05 ago, expires in 00:42:52, sent 02:37:54 ago
Router    *0.0.0.1          10.255.4.93      0x80000002  1096  0xc677  40
    bits 0x3, Options 0x13
    Type PointToPoint (1), Metric 10
    Loc-If-Id 3, Nbr-If-Id 2, Nbr-Rtr-Id 10.255.4.85
    Gen timer 00:00:40
    Aging timer 00:41:44
    Installed 00:18:16 ago, expires in 00:41:44, sent 00:18:14 ago
    Ours
InterArPfx *0.0.0.2          10.255.4.93      0x80000001  1053  0xb96f  44
    Prefix feee::10:10:2:0/126
    Prefix-options 0x0, Metric 10
    Gen timer 00:17:02
    Aging timer 00:42:26
    Installed 00:17:33 ago, expires in 00:42:27, sent 00:17:31 ago
    Ours
InterArPfx *0.0.0.3          10.255.4.93      0x80000001  1053  0x71d3  44
    Prefix feee::10:255:4:97/128
    Prefix-options 0x0, Metric 10
    Gen timer 00:21:07
    Aging timer 00:42:26
    Installed 00:17:33 ago, expires in 00:42:27, sent 00:17:31 ago
    Ours
InterArRtr *0.0.0.1          10.255.4.93      0x80000001  1053  0xe159  32
    Dest-router-id 10.255.4.97, Options 0x19, Metric 10
    Gen timer 00:29:18
    Aging timer 00:42:26
    Installed 00:17:33 ago, expires in 00:42:27, sent 00:17:31 ago
    Ours
IntraArPfx 0.0.0.1          10.255.4.85      0x80000002  1028  0x2403  72
    Ref-lsa-type Router, Ref-lsa-id 0.0.0.0, Ref-router-id 10.255.4.85
    Prefix-count 2
    Prefix feee::10:255:4:85/128
    Prefix-options 0x2, Metric 0
    Prefix feee::10:10:1:0/126
    Prefix-options 0x0, Metric 10
    Aging timer 00:42:51
    Installed 00:17:05 ago, expires in 00:42:52, sent 02:37:54 ago
IntraArPfx *0.0.0.1          10.255.4.93      0x80000002  575  0x788f  72
    Ref-lsa-type Router, Ref-lsa-id 0.0.0.0, Ref-router-id 10.255.4.93
    Prefix-count 2
    Prefix feee::10:255:4:93/128
    Prefix-options 0x2, Metric 0
    Prefix feee::10:10:1:0/126
    Prefix-options 0x0, Metric 10
    Gen timer 00:33:23
    Aging timer 00:50:24
    Installed 00:09:35 ago, expires in 00:50:25, sent 00:09:33 ago
    OSPF3 link state database, area 0.0.0.1
Type      ID              Adv Rtr          Seq              Age  Cksum  Len
Router    *0.0.0.1          10.255.4.93      0x80000003  1059  0xea40  40
    bits 0x3, Options 0x19
    Type Transit (2), Metric 10
    Loc-If-Id 2, Nbr-If-Id 2, Nbr-Rtr-Id 10.255.4.97
    Gen timer 00:08:51
    Aging timer 00:42:20
    Installed 00:17:39 ago, expires in 00:42:21, sent 00:17:37 ago
Router     0.0.0.1          10.255.4.97      0x80000006  994  0xc95b  40
    bits 0x2, Options 0x19

```



```

Type Transit (2), Metric 10
  Loc-If-Id 2, Nbr-If-Id 2, Nbr-Rtr-Id 10.255.4.97
Aging timer 00:43:25
  Installed 00:16:31 ago, expires in 00:43:26, sent 02:37:54 ago
Network    0.0.0.2          10.255.4.97      0x80000002  1059  0x4598  32
Options 0x11
  Attached router 10.255.4.97
  Attached router 10.255.4.93
Aging timer 00:42:20
  Installed 00:17:36 ago, expires in 00:42:21, sent 02:37:54 ago
InterArPfx *0.0.0.1      10.255.4.93      0x80000002   260  0xa980  44
  Prefix feee::10:10:1:0/126
  Prefix-options 0x0, Metric 10
  Gen timer 00:45:39
  Aging timer 00:55:39
  Installed 00:04:20 ago, expires in 00:55:40, sent 00:04:18 ago
  Ours
InterArPfx *0.0.0.2      10.255.4.93      0x80000002   205  0xd47e  44
  Prefix feee::10:255:4:93/128
  Prefix-options 0x0, Metric 0
  Gen timer 00:46:35
  Aging timer 00:56:35
  Installed 00:03:25 ago, expires in 00:56:35, sent 00:03:23 ago
  Ours
InterArPfx *0.0.0.3      10.255.4.93      0x80000001  1089  0x9bbb  44
  Prefix feee::10:255:4:85/128
  Prefix-options 0x0, Metric 10
  Gen timer 00:04:46
  Aging timer 00:41:51
  Installed 00:18:09 ago, expires in 00:41:51, sent 00:17:43 ago
  Ours
NSSA      0.0.0.1          10.255.4.97      0x80000002   505  0x45ee  44
  Prefix feee::200:200:1:0/124
  Prefix-options 0x8, Metric 10, Type 2,
  Aging timer 00:51:35
  Installed 00:08:22 ago, expires in 00:51:35, sent 02:37:54 ago
IntraArPfx 0.0.0.1      10.255.4.97      0x80000006   994  0x2f77  52
  Ref-lsa-type Router, Ref-lsa-id 0.0.0.0, Ref-router-id 10.255.4.97
  Prefix-count 1
  Prefix feee::10:255:4:97/128
  Prefix-options 0x2, Metric 0
  Aging timer 00:43:25
  Installed 00:16:31 ago, expires in 00:43:26, sent 02:37:54 ago
IntraArPfx 0.0.0.3      10.255.4.97      0x80000002  1059  0x4446  52
  Ref-lsa-type Network, Ref-lsa-id 0.0.0.2, Ref-router-id 10.255.4.97
  Prefix-count 1
  Prefix feee::10:10:2:0/126
  Prefix-options 0x0, Metric 0
  Aging timer 00:42:20
  Installed 00:17:36 ago, expires in 00:42:21, sent 02:37:54 ago
  OSPF3 AS SCOPE link state database
Type      ID          Adv Rtr          Seq          Age  Cksum  Len
Extern    0.0.0.1          10.255.4.85      0x80000002   206  0x9b86  44
  Prefix feee::100:100:1:0/124
  Prefix-options 0x0, Metric 20, Type 2,
  Aging timer 00:56:34
  Installed 00:03:23 ago, expires in 00:56:34, sent 02:37:54 ago
Extern    *0.0.0.1          10.255.4.93      0x80000001  1053  0x59c9  44
  Prefix feee::200:200:1:0/124
  Prefix-options 0x0, Metric 10, Type 2,
  Gen timer 00:25:12

```

Aging timer 00:42:26

Installed 00:17:33 ago, expires in 00:42:27, sent 00:17:31 ago

OSPF3 Link-Local link state database, interface ge-1/3/0.0

Type	ID	Adv Rtr	Seq	Age	Cksum	Len
Link	*0.0.0.2	10.255.4.93	0x80000003	1059	0x4dab	64

fe80::290:69ff:fe39:1cdb

Options 0x11, priority 128

Prefix-count 1

Prefix feee::10:10:2:0/126 Prefix-options 0x0

Gen timer 00:12:56

Aging timer 00:42:20

Installed 00:17:39 ago, expires in 00:42:21, sent 00:17:37 ago

Link	0.0.0.2	10.255.4.97	0x80000003	205	0xa87d	64
------	---------	-------------	------------	-----	--------	----

fe80::290:69ff:fe38:883e

Options 0x11, priority 128

Prefix-count 1

Prefix feee::10:10:2:0/126 Prefix-options 0x0

Aging timer 00:56:35

Installed 00:03:22 ago, expires in 00:56:35, sent 02:37:54 ago

OSPF3 Link-Local link state database, interface so-2/2/0.0

Type	ID	Adv Rtr	Seq	Age	Cksum	Len
Link	0.0.0.2	10.255.4.85	0x80000002	506	0x42bb	64

fe80::280:42ff:fe10:f169

Options 0x13, priority 128

Prefix-count 1

Prefix feee::10:10:1:0/126 Prefix-options 0x0

Aging timer 00:51:34

Installed 00:08:23 ago, expires in 00:51:34, sent 02:37:54 ago

Link	*0.0.0.3	10.255.4.93	0x80000002	505	0x6b7a	64
------	----------	-------------	------------	-----	--------	----

fe80::280:42ff:fe10:f177

Options 0x13, priority 128

Prefix-count 1

Prefix feee::10:10:1:0/126 Prefix-options 0x0

Gen timer 00:37:28

Aging timer 00:51:35

Installed 00:08:25 ago, expires in 00:51:35, sent 00:08:23 ago

Ours

show ospf3 database summary

user@host> show ospf3 database summary

Area 0.0.0.0:

2 Router LSAs

1 InterArPfx LSAs

1 InterArRtr LSAs

1 IntraArPfx LSAs

Area 0.0.0.1:

2 Router LSAs

1 Network LSAs

2 InterArPfx LSAs

1 NSSA LSAs

1 IntraArPfx LSAs

Externals:

2 Extern LSAs

Interface ge-1/3/0.0:

1 Link LSAs

Interface lo0.0:

Interface so-2/2/0.0:

1 Link LSAs

show (ospf | ospf3) interface

Syntax	<code>show (ospf ospf3) interface</code> <code><brief detail extensive></code> <code><area <i>area-id</i>></code> <code><interface-name></code> <code><instance <i>instance-name</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code> <code><realm (ipv4-multicast ipv4-unicast ipv6-multicast)></code>
Syntax (EX Series Switches and QFX Series)	<code>show (ospf ospf3) interface</code> <code><brief detail extensive></code> <code><area <i>area-id</i>></code> <code><interface-name></code> <code><instance <i>instance-name</i>></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. area option introduced in Junos OS Release 9.2. area option introduced in Junos OS Release 9.2 for EX Series switches. realm option introduced in Junos OS Release 9.2. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display the status of OSPF interfaces.
Options	none —Display standard information about the status of all OSPF interfaces for all routing instances brief detail extensive —(Optional) Display the specified level of output. area <i>area-id</i> —(Optional) Display information about the interfaces that belong to the specified area. <i>interface-name</i> —(Optional) Display information for the specified interface. instance <i>instance-name</i> —(Optional) Display all OSPF interfaces under the named routing instance. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. realm (ipv4-multicast ipv4-unicast ipv6-multicast) —(OSPFv3 only) (Optional) Display information about the interfaces for the specified OSPFv3 realm, or address family. Use the realm option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.
Required Privilege Level	view
List of Sample Output	show ospf interface brief on page 733 show ospf interface detail on page 733 show ospf3 interface detail on page 733

[show ospf interface detail \(When Multiarea Adjacency Is Configured\) on page 733](#)

[show ospf interface area area-id on page 734](#)

[show ospf interface extensive \(When Flooding Reduction Is Enabled\) on page 735](#)

[show ospf interface extensive \(When LDP Synchronization Is Configured\) on page 735](#)

Output Fields Table 113 on page 731 lists the output fields for the **show (ospf | ospf3) interface** command. Output fields are listed in the approximate order in which they appear.

Table 113: show (ospf | ospf3) interface Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface running OSPF version 2 or OSPF version 3.	All levels
State	State of the interface: BDR , Down , DR , DRother , Loop , PtToPt , or Waiting .	All levels
Area	Number of the area that the interface is in.	All levels
DR ID	Address of the area's designated router.	All levels
BDR ID	Backup designated router for a particular subnet.	All levels
Nbrs	Number of neighbors on this interface.	All levels
Type	Type of interface: LAN , NBMA , P2MP , P2P , or Virtual .	detail extensive
Address	IP address of the neighbor.	detail extensive
Mask	Netmask of the neighbor.	detail extensive
Prefix-length	(OSPFv3) IPv6 prefix length, in bits.	detail extensive
OSPF3-Intf-Index	(OSPFv3) OSPF version 3 interface index.	detail extensive
MTU	Interface maximum transmission unit (MTU).	detail extensive
Cost	Interface cost (metric).	detail extensive
DR addr	Address of the designated router.	detail extensive
BDR addr	Address of the backup designated router.	detail extensive
Adj count	Number of adjacent neighbors.	detail extensive
Secondary	Indicates that this interface is configured as a secondary interface for this area. This interface can belong to more than one area, but can be designated as a primary interface for only one area.	detail extensive
Flood Reduction	Indicates that this interface is configured with flooding reduction. All self-originated LSAs from this interface are initially sent with the DoNotAge bit set. As a result, LSAs are refreshed only when a change occurs.	extensive

Table 113: show (ospf | ospf3) interface Output Fields (*continued*)

Field Name	Field Description	Level of Output
Priority	Router priority used in designated router (DR) election on this interface.	detail extensive
Flood list	List of link-state advertisements (LSAs) that might be about to flood this interface.	extensive
Ack list	Acknowledgment list. List of pending acknowledgments on this interface.	extensive
Descriptor list	List of packet descriptors.	extensive
Hello	Configured value for the hello timer.	detail extensive
Dead	Configured value for the dead timer.	detail extensive
Auth type	(OSPFv2) Authentication mechanism for sending and receiving OSPF protocol packets: <ul style="list-style-type: none"> • MD5—The MD5 mechanism is configured in accordance with RFC 2328. • None—No authentication method is configured. • Password—A simple password (RFC 2328) is configured. 	detail extensive
Topology	(Multiarea adjacency) Name of topology: default or name .	
LDP sync state	(OSPFv2 and LDP synchronization) Current state of LDP synchronization: in sync , in holddown , and not supported .	extensive
reason	(OSPFv2 and LDP synchronization) Reason for the current state of LDP synchronization. The LDP session might be up or down, or adjacency might be up or down.	extensive
config holdtime	(OSPFv2 and LDP synchronization) Configured value of the hold timer. If the state is not synchronized, and the hold time is not infinity, the remaining field displays the number of seconds that remain until the configured hold timer expires.	extensive
IPSec SA name	(OSPFv2) Name of the IPSec security association name.	detail extensive
Active key ID	(OSPFv2 and MD5) Number from 0 to 255 that uniquely identifies an MD5 key.	detail extensive
Start time	(OSPFv2 and MD5) Time at which the routing device starts using an MD5 key to authenticate OSPF packets transmitted on the interface on which this key is configured. To authenticate received OSPF protocol packets, the key becomes effective immediately after the configuration is committed. If the start time option is not configured, the key is effective immediately for send and receive and is displayed as Start time 1970 Jan 01 00:00:00 PST .	detail extensive
ReXmit	Configured value for the Retransmit timer.	detail extensive
Stub, Not Stub, or Stub NSSA	Type of area.	detail extensive

Sample Output

show ospf interface brief

```
user@host> show ospf interface brief
```

Intf	State	Area	DR ID	BDR ID	Nbrs
at-5/1/0.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
ge-2/3/0.0	DR	0.0.0.0	192.168.4.16	192.168.4.15	1
lo0.0	DR	0.0.0.0	192.168.4.16	0.0.0.0	0
so-0/0/0.0	Down	0.0.0.0	0.0.0.0	0.0.0.0	0
so-6/0/1.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1
so-6/0/2.0	Down	0.0.0.0	0.0.0.0	0.0.0.0	0
so-6/0/3.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1

show ospf interface detail

```
user@host> show ospf interface detail
```

Interface	State	Area	DR ID	BDR ID	Nbrs
fe-0/0/1.0	BDR	0.0.0.0	192.168.37.12	10.255.245.215	1

Type LAN, address 192.168.37.11, Mask 255.255.255.248, MTU 4460, Cost 40
 DR addr 192.168.37.12, BDR addr 192.168.37.11, Adj count 1, Priority 128
 Hello 10, Dead 40, ReXmit 5, Not Stub

Interface	State	Area	DR ID	BDR ID	Nbrs
tl-0/2/1.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	0

Type P2P, Address 0.0.0.0, Mask 0.0.0.0, MTU 1500, Cost 2604
 Adj count 0
 Hello 10, Dead 40, ReXmit 5, Not Stub
 Auth type: MD5, Active key ID 3, Start time 2002 Nov 19 10:00:00 PST
 IPsec SA Name: sa

show ospf3 interface detail

```
user@host> show ospf3 interface so-0/0/3.0 detail
```

Interface	State	Area	DR-ID	BDR-ID	Nbrs
so-0/0/3.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1

Address fe80::2a0:a5ff:fe28:1dfc, Prefix-length 64
 OSPF3-Intf-index 1, Type P2P, MTU 4470, Cost 12, Adj-count 1
 Hello 10, Dead 40, ReXmit 5, Not Stub

show ospf interface detail (When Multiarea Adjacency Is Configured)

```
user@host> show ospf interface detail
```

```
regress@router> show ospf interface detail
```

Interface	State	Area	DR ID	BDR ID	Nbrs
lo0.0	DR	0.0.0.0	10.255.245.2	0.0.0.0	0

Type: LAN, Address: 127.0.0.1, Mask: 255.255.255.255, MTU: 65535, Cost: 0
 DR addr: 127.0.0.1, Adj count: 0, Priority: 128
 Hello: 10, Dead: 40, ReXmit: 5, Not Stub
 Auth type: None
 Topology default (ID 0) -> Cost: 0

Interface	State	Area	DR ID	BDR ID	Nbrs
lo0.0	DR	0.0.0.0	10.255.245.2	0.0.0.0	0

Type: LAN, Address: 10.255.245.2, Mask: 255.255.255.255, MTU: 65535, Cost: 0
 DR addr: 10.255.245.2, Adj count: 0, Priority: 128
 Hello: 10, Dead: 40, ReXmit: 5, Not Stub
 Auth type: None
 Topology default (ID 0) -> Cost: 0

Interface	State	Area	DR ID	BDR ID	Nbrs
so-0/0/0.0	PtToPt	0.0.0.0	0.0.0.0	0.0.0.0	1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1

```

Adj count: 1
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-0/0/0.0      PtToPt  0.0.0.0      0.0.0.0      0.0.0.0      0

Type: P2P, Address: 192.168.37.46, Mask: 255.255.255.254, MTU: 4470, Cost: 1
Adj count: 0, , Passive
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Passive, Cost: 1
so-1/0/0.0      PtToPt  0.0.0.0      0.0.0.0      0.0.0.0      1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-1/0/0.0      PtToPt  0.0.0.0      0.0.0.0      0.0.0.0      0

Type: P2P, Address: 192.168.37.54, Mask: 255.255.255.254, MTU: 4470, Cost: 1
Adj count: 0, , Passive
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Passive, Cost: 1
so-0/0/0.0      PtToPt  1.1.1.1      0.0.0.0      0.0.0.0      1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-1/0/0.0      PtToPt  1.1.1.1      0.0.0.0      0.0.0.0      1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-0/0/0.0      PtToPt  2.2.2.2      0.0.0.0      0.0.0.0      1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
so-1/0/0.0      PtToPt  2.2.2.2      0.0.0.0      0.0.0.0      1

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
Adj count: 1, Secondary
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1

```

show ospf interface area area-id

```

user@host> show ospf interface area 1.1.1.1
Interface      State   Area      DR ID      BDR ID      Nbrs
so-0/0/0.0     PtToPt  1.1.1.1   0.0.0.0    0.0.0.0     1
so-1/0/0.0     PtToPt  1.1.1.1   0.0.0.0    0.0.0.0     1

```


show ospf interface extensive (When Flooding Reduction Is Enabled)

```
user@host> show ospf interface extensive
Interface          State   Area      DR ID      BDR ID      Nbrs
fe-0/0/0.0         PtToPt 0.0.0.0    0.0.0.0    0.0.0.0      0

Type: P2P, Address: 10.10.10.1, Mask: 255.255.255.0, MTU: 1500, Cost: 1
Adj count: 0
Secondary, Flood Reduction
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
Topology default (ID 0) -> Cost: 1
```

show ospf interface extensive (When LDP Synchronization Is Configured)

```
user@host> show ospf interface extensive
Interface          State   Area      DR ID      BDR ID
Nbrs
so-1/0/3.0         Down    0.0.0.0    0.0.0.0    0.0.0.0
0

Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 65535
Adj count: 0
Hello: 10, Dead: 40, ReXmit: 5, Not Stub
Auth type: None
LDP sync state: in holddown, for: 00:00:08, reason: LDP down during config
config holddown: 10 seconds, remaining: 1
```

show (ospf | ospf3) io-statistics

Syntax	show (ospf ospf3) io-statistics <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch and QFX Series)	show (ospf ospf3) io-statistics
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display Open Shortest Path First (OSPF) input and output statistics.
Options	<p>none—Display OSPF input and output statistics.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> clear (ospf ospf3) statistics on page 693
List of Sample Output	show ospf io-statistics on page 736
Output Fields	<p>Table 114 on page 736 lists the output fields for the show ospf io-statistics command. Output fields are listed in the approximate order in which they appear.</p>

Table 114: show (ospf | ospf3) io-statistics Output Fields

Field Name	Field Description
Packets read	Number of OSPF packets read since the last time the routing protocol was started.
average per run	Total number of packets divided by the total number of times the OSPF read operation is scheduled to run.
max run	Maximum number of packets for a given run among all scheduled runs.
Receive errors	Number of faulty packets received with errors.

Sample Output

show ospf io-statistics

```
user@host> show ospf io-statistics
```

```
Packets read: 7361, average per run: 1.00, max run: 1
```

Receive errors:
None

show (ospf | ospf3) log

Syntax	<pre>show (ospf ospf3) log <instance <i>instance-name</i>> <logical-system (all <i>logical-system-name</i>)> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)> <topology <i>topology-name</i>></pre>
Syntax (EX Series Switch and QFX Series)	<pre>show (ospf ospf3) log <instance <i>instance-name</i>> <topology <i>topology-name</i>></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>topology option introduced in Junos OS Release 9.0.</p> <p>topology option introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>realm option introduced in Junos OS Release 9.2.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Display the entries in the Open Shortest Path First (OSPF) log of SPF calculations.
Options	<p>none—Display entries in the OSPF log of SPF calculations for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Display entries for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>topology <i>topology-name</i>—(Optional) (OSPFv2 only) Display entries for the specified topology.</p> <p>realm (ipv4-multicast ipv4-unicast ipv6-multicast)—(OSPFv3 only) (Optional) Display entries for the specified OSPFv3 realm, or address family. Use the realm option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p>
Required Privilege Level	view
List of Sample Output	<p>show ospf log on page 739</p> <p>show ospf log topology voice on page 739</p>
Output Fields	<p>Table 115 on page 738 lists the output fields for the show (ospf ospf3) log command. Output fields are listed in the approximate order in which they appear.</p>

Table 115: show (ospf | ospf3) log Output Fields

Field Name	Field Description
When	Time, in weeks (w) and days (d), since the SPF calculation was made.

Table 115: show (ospf | ospf3) log Output Fields (*continued*)

Field Name	Field Description
Type	Type of calculation: Cleanup, External, Interarea, NSSA, Redist, SPF, Stub, Total, or VirtualLink.
Elapsed	Amount of time, in seconds, that elapsed during the operation, or the time required to complete the SPF calculation. The start time is the time displayed in the When field.

Sample Output

show ospf log

```

user@host> show ospf log
When          Type          Elapsed
1w4d 17:25:58 Stub          0.000017
1w4d 17:25:58 SPF           0.000070
1w4d 17:25:58 Stub          0.000019
1w4d 17:25:58 Interarea     0.000054
1w4d 17:25:58 External      0.000005
1w4d 17:25:58 Cleanup       0.000203
1w4d 17:25:58 Total         0.000537
1w4d 17:24:48 SPF           0.000125
1w4d 17:24:48 Stub          0.000017
1w4d 17:24:48 SPF           0.000100
1w4d 17:24:48 Stub          0.000016
1w4d 17:24:48 Interarea     0.000056
1w4d 17:24:48 External      0.000005
1w4d 17:24:48 Cleanup       0.000238
1w4d 17:24:48 Total         0.000600
...

```

show ospf log topology voice

```

user@host> show ospf log topology voice
Topology voice SPF log:

    Last instance of each event type
When          Type          Elapsed
00:06:11      SPF           0.000116
00:06:11      Stub          0.000114
00:06:11      Interarea     0.000126
00:06:11      External      0.000067
00:06:11      NSSA          0.000037
00:06:11      Cleanup       0.000186

    Maximum length of each event type
When          Type          Elapsed
00:13:43      SPF           0.000140
00:13:33      Stub          0.000116
00:13:43      Interarea     0.000128
00:13:33      External      0.000075
00:13:38      NSSA          0.000039
00:13:53      Cleanup       0.000657

    Last 100 events

```

When	Type	Elapsed
00:13:53	SPF	0.000090
00:13:53	Stub	0.000041
00:13:53	Interarea	0.000123
00:13:53	External	0.000040
00:13:53	NSSA	0.000038
00:13:53	Cleanup	0.000657
00:13:53	Total	0.001252
.		
.		
00:06:11	SPF	0.000116
00:06:11	Stub	0.000114
00:06:11	Interarea	0.000126
00:06:11	External	0.000067
00:06:11	NSSA	0.000037
00:06:11	Cleanup	0.000186
00:06:11	Total	0.000818

show (ospf | ospf3) neighbor

Syntax	<pre>show (ospf ospf3) neighbor <brief detail extensive> <area <i>area-id</i>> <instance (all <i>instance-name</i>)> <interface <i>interface-name</i>> <logical-system (all <i>logical-system-name</i>)> <neighbor> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show (ospf ospf3) neighbor <brief detail extensive> <area <i>area-id</i>> <instance (all <i>instance-name</i>)> <interface <i>interface-name</i>> <neighbor></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>instance all option introduced in Junos OS Release 9.1.</p> <p>instance all option introduced in Junos OS Release 9.1 for EX Series switches.</p> <p>area, interface, and realm options introduced in Junos OS Release 9.2.</p> <p>area and interface options introduced in Junos OS Release 9.2 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	<p>Display information about OSPF neighbors.</p> <p>CPU utilization might increase while the device learns its OSPF neighbors. We recommend that you use the show (ospf ospf3) neighbor command after the device learns and establishes OSPF neighbor adjacencies. Depending on the size of your network, this might take several minutes. If you receive a “timeout communicating with routing daemon” error when using the show (ospf ospf3) neighbor command, wait several minutes before attempting to use the command again. This is not a critical system error, but you might experience a delay in using the CLI.</p>
Options	<p>none—Display standard information about all OSPF neighbors for all routing instances.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>area <i>area-id</i>—(Optional) Display information about the OSPF neighbors for the specified area.</p> <p>instance (all <i>instance-name</i>)—(Optional) Display all OSPF interfaces for all routing instances or under the named routing instance.</p> <p>interface <i>interface-name</i>—(Optional) Display information about OSPF neighbors for the specified logical interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>

neighbor—(Optional) Display information about the specified OSPF neighbor.

realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)—(OSPFv3 only) (Optional) Display information about the OSPF neighbors for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

Required Privilege Level view

Related Documentation

- [clear \(ospf | ospf3\) neighbor on page 690](#)

List of Sample Output

- [show ospf neighbor brief on page 744](#)
- [show ospf neighbor detail on page 744](#)
- [show ospf neighbor extensive on page 745](#)
- [show ospf3 neighbor detail on page 746](#)
- [show ospf neighbor area area-id on page 746](#)
- [show ospf neighbor interface interface-name on page 746](#)
- [show ospf3 neighbor instance all \(OSPFv3 Multiple Family Address Support Enabled\) on page 746](#)

Output Fields [Table 116 on page 742](#) lists the output fields for the **show (ospf | ospf3) neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 116: show (ospf | ospf3) neighbor Output Fields

Field Name	Field Description	Level of Output
Address	Address of the neighbor.	All levels
Interface	Interface through which the neighbor is reachable.	All levels

Table 116: show (ospf | ospf3) neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
State	<p>State of the neighbor:</p> <ul style="list-style-type: none"> • Attempt—Valid only for neighbors attached to nonbroadcast networks. It indicates that no recent information has been received from the neighbor, but that a more concerted effort must be made to contact the neighbor. • Down—Initial state of a neighbor conversation. It indicates that no recent information has been received from the neighbor. Hello packets might continue to be sent to neighbors in the Down state, although at a reduced frequency. • Exchange—Routing device is describing its entire link-state database by sending database description packets to the neighbor. Each packet has a sequence number and is explicitly acknowledged. • ExStart—First step in creating an adjacency between the two neighboring routing devices. The goal of this step is to determine which routing device is the master, and to determine the initial sequence number. • Full—Neighboring routing devices are fully adjacent. These adjacencies appear in router link and network link advertisements. • Init—A hello packet has recently been sent by the neighbor. However, bidirectional communication has not yet been established with the neighbor. This state might occur, for example, because the routing device itself did not appear in the neighbor's hello packet. • Loading—Link-state request packets are sent to the neighbor to acquire more recent advertisements that have been discovered (but not yet received) in the Exchange state. • 2Way—Communication between the two routing devices is bidirectional. This state has been ensured by the operation of the Hello Protocol. This is the most advanced state short of beginning adjacency establishment. The (backup) designated router is selected from the set of neighbors in state 2Way or greater. 	All levels
ID	Router ID of the neighbor.	All levels
Pri	Priority of the neighbor to become the designated router.	All levels
Dead	Number of seconds until the neighbor becomes unreachable.	All levels
Link state acknowledgment list	Number of link-state acknowledgments received.	extensive
Link state retransmission list	<p>Total number of link-state advertisements retransmitted. For extensive output only, the following information is also displayed:</p> <ul style="list-style-type: none"> • Type—Type of link advertisement: ASBR, Sum, Extern, Network, NSSA, OpaqArea, Router, or Summary. • LSA ID—LSA identifier included in the advertisement. An asterisk preceding the identifier marks database entries that originated from the local routing device. • Adv rtr—Address of the routing device that sent the advertisement. • Seq—Link sequence number of the advertisement. 	detail extensive

Table 116: show (ospf | ospf3) neighbor Output Fields (*continued*)

Field Name	Field Description	Level of Output
Neighbor-address	(OSPFv3 only) If the neighbor uses virtual links, the Neighbor-address is the site-local, local, or global address. If the neighbor uses a physical interface, the Neighbor-address is an IPv6 link-local address.	detail extensive
area	Area that the neighbor is in.	detail extensive
OSPF3-Intf-Index	(OSPFv3 only) Displays the OSPFv3 interface index.	detail extensive
opt	Option bits received in the hello packets from the neighbor.	detail extensive
DR or DR-ID	Address of the designated router.	detail extensive
BDR or BDR-ID	Address of the backup designated router.	detail extensive
Up	Length of time since the neighbor came up.	detail extensive
adjacent	Length of time since the adjacency with the neighbor was established.	detail extensive

Sample Output

show ospf neighbor brief

```

user@host> show ospf neighbor brief
  Address      Intf      State      ID          Pri  Dead
192.168.254.225 fxp3.0    2Way       10.250.240.32 128  36
192.168.254.230 fxp3.0    Full       10.250.240.8  128  38
192.168.254.229 fxp3.0    Full       10.250.240.35 128  33
10.1.1.129      fxp2.0    Full       10.250.240.12 128  37
10.1.1.131      fxp2.0    Full       10.250.240.11 128  38
10.1.2.1        fxp1.0    Full       10.250.240.9  128  32
10.1.2.81       fxp0.0    Full       10.250.240.10 128  33

```

show ospf neighbor detail

```

user@host> show ospf neighbor detail
  Address      Interface      State      ID          Pri  Dead
10.5.1.2      ge-1/2/0.1     Full       10.5.1.2    128  37
area 0.0.0.1, opt 0x42, DR 10.5.1.2, BDR 10.5.1.1
Up 06:09:28, adjacent 05:17:36
Link state acknowledgment list: 3 entries

Link state retransmission list: 9 entries

10.5.10.2     ge-1/2/0.10    ExStart    10.5.1.38   128  34
area 0.0.0.1, opt 0x42, DR 10.5.10.2, BDR 10.5.10.1
Up 06:09:28
master, seq 0xac1530f8, rexmit DBD in 3 sec
rexmit LSREQ in 0 sec
10.5.11.2     ge-1/2/0.11    Full       10.5.1.42   128  38
area 0.0.0.1, opt 0x42, DR 10.5.11.2, BDR 10.5.11.1
Up 06:09:28, adjacent 05:26:46
Link state retransmission list: 1 entries

```

```

10.5.12.2      ge-1/2/0.12      ExStart  10.5.1.46      128    33
area 0.0.0.1, opt 0x42, DR 10.5.12.2, BDR 10.5.12.1
Up 06:09:28
master, seq 0xac188a68, rexmit DBD in 2 sec
rexmit LSREQ in 0 sec

```

show ospf neighbor extensive

```

user@host> show ospf neighbor extensive
Address      Interface      State      ID      Pri  Dead
10.5.1.2      ge-1/2/0.1     Full       10.5.1.2 128   33
area 0.0.0.1, opt 0x42, DR 10.5.1.2, BDR 10.5.1.1
Up 06:09:42, adjacent 05:17:50
Link state retransmission list:

  Type      LSA ID      Adv rtr      Seq
  Summary   10.8.56.0    172.25.27.82 0x8000004d
  Router    10.5.1.94    10.5.1.94    0x8000005c
  Network   10.5.24.2    10.5.1.94    0x80000036
  Summary   10.8.57.0    172.25.27.82 0x80000024
  Extern    1.10.90.0    10.8.1.2     0x80000041
  Extern    1.4.109.0    10.6.1.2     0x80000041
  Router    10.5.1.190   10.5.1.190   0x8000005f
  Network   10.5.48.2    10.5.1.190   0x8000003d
  Summary   10.8.58.0    172.25.27.82 0x8000004d
  Extern    1.10.91.0    10.8.1.2     0x80000041
  Extern    1.4.110.0    10.6.1.2     0x80000041
  Router    10.5.1.18    10.5.1.18    0x8000005f
  Network   10.5.5.2     10.5.1.18    0x80000033
  Summary   10.8.59.0    172.25.27.82 0x8000003a
  Summary   10.8.62.0    172.25.27.82 0x80000025

10.5.10.2     ge-1/2/0.10     ExStart  10.5.1.38      128    38
area 0.0.0.1, opt 0x42, DR 10.5.10.2, BDR 10.5.10.1
Up 06:09:42
master, seq 0xac1530f8, rexmit DBD in 2 sec
rexmit LSREQ in 0 sec
10.5.11.2     ge-1/2/0.11     Full     10.5.1.42      128    33
area 0.0.0.1, opt 0x42, DR 10.5.11.2, BDR 10.5.11.1
Up 06:09:42, adjacent 05:27:00
Link state retransmission list:

  Type      LSA ID      Adv rtr      Seq
  Summary   10.8.58.0    172.25.27.82 0x8000004d

```

Extern	1.10.91.0	10.8.1.2	0x80000041
Extern	1.1.247.0	10.5.1.2	0x8000003f
Extern	1.4.110.0	10.6.1.2	0x80000041
Router	10.5.1.18	10.5.1.18	0x8000005f
Network	10.5.5.2	10.5.1.18	0x80000033
Summary	10.8.59.0	172.25.27.82	0x8000003a

show ospf3 neighbor detail

```

user@host> show ospf3 neighbor detail
ID          Interface          State    Pri    Dead
10.255.71.13 fe-0/0/2.0          Full     128    30
Neighbor-address fe80::290:69ff:fe9b:e002
area 0.0.0.0, opt 0x13, OSPF3-Intf-Index 2
DR-ID 10.255.71.13, BDR-ID 10.255.71.12
Up 02:51:43, adjacent 02:51:43

```

show ospf neighbor area area-id

```

user@host >show ospf neighbor area 1.1.1.1
Address      Interface          State    ID          Pri    Dead
192.168.37.47 so-0/0/0.0        Full     10.255.245.4 128    33
Area 1.1.1.1
192.168.37.55 so-1/0/0.0        Full     10.255.245.5 128    37
Area 1.1.1.1

```

show ospf neighbor interface interface-name

```

user@host >show ospf neighbor interface so-0/0/0.0
Address      Interface          State    ID          Pri    Dead
192.168.37.47 so-0/0/0.0        Full     10.255.245.4 128    37
Area 0.0.0.0
192.168.37.47 so-0/0/0.0        Full     10.255.245.4 128    33
Area 1.1.1.1
192.168.37.47 so-0/0/0.0        Full     10.255.245.4 128    32
Area 2.2.2.2

```

show ospf3 neighbor instance all (OSPFv3 Multiple Family Address Support Enabled)

```

user @host > show ospf3 neighbor instance all
Instance: ina
Realm: ipv6-unicast
ID          Interface          State    Pri    Dead
100.1.1.1    fe-0/0/2.0          Full     128    37
Neighbor-address fe80::217:cb00:c87c:8c03
Instance: inb
Realm: ipv4-unicast
ID          Interface          State    Pri    Dead
100.1.2.1    fe-0/0/2.1          Full     128    33
Neighbor-address fe80::217:cb00:c97c:8c03

```

show (ospf | ospf3) overview

Syntax	show (ospf ospf3) overview <brief extensive> <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)>
Syntax (EX Series Switch and QFX Series)	show (ospf ospf3) overview <brief extensive> <instance <i>instance-name</i> >
Release Information	Command introduced in Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. realm option introduced in Junos OS Release 9.2. Database protection introduced in Junos 10.2. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display Open Shortest Path First (OSPF) overview information.
Options	<p>none—Display standard information about all OSPF neighbors for all routing instances.</p> <p>brief extensive—(Optional) Display the specified level of output.</p> <p>instance <i>instance-name</i>—(Optional) Display all OSPF interfaces under the named routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>realm (ipv4-multicast ipv4-unicast ipv6-multicast)—(Optional) (OSPFv3 only) Display information about the specified OSPFv3 realm, or address family. Use the realm option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p>
Required Privilege Level	view
List of Sample Output	show ospf overview on page 749 show ospf overview (With Database Protection) on page 750 show ospf3 overview (With Database Protection) on page 750 show ospf overview extensive on page 750
Output Fields	Table 117 on page 747 lists the output fields for the show ospf overview command. Output fields are listed in the approximate order in which they appear.

Table 117: show ospf overview Output Fields

Field name	Field Description	Level of Output
Instance	OSPF routing instance.	All levels

Table 117: show ospf overview Output Fields (*continued*)

Field name	Field Description	Level of Output
Router ID	Router ID of the routing device.	All levels
Route table index	Route table index.	All levels
Configured overload	Overload capability is enabled. If the overload timer is also configured, display the time that remains before it is set to expire. This field is not displayed after the timer expires.	All levels
Topology	Topology identifier.	All levels
Prefix export count	Number of prefixes exported into OSPF.	All levels
Full SPF runs	Number of complete Shortest Path First calculations.	All levels
SPF delay	Delay before performing consecutive Shortest Path First calculations.	All levels
SPF holddown	Delay before performing additional Shortest Path First (SPF) calculations after the maximum number of consecutive SPF calculations is reached.	All levels
SPF rapid runs	Maximum number of Shortest Path First calculations that can be performed in succession before the hold-down timer begins.	All levels
LSA refresh time	Refresh period for link-state advertisement (in minutes).	All levels
Database protection state	Current state of database protection.	All levels
Warning threshold	Threshold at which a warning message is logged (percentage of maximum LSA count).	All levels
Non self-generated LSAs	Number of LSAs whose router ID is not equal to the local router ID: Current , Warning (threshold), and Allowed .	All levels
Ignore time	How long the database has been in the ignore state.	All levels
Reset time	How long the database must stay out of the ignore or isolated state before it returns to normal operations.	All levels
Ignore count	Number of times the database has been in the ignore state: Current and Allowed .	All levels
Restart	Graceful restart capability: enabled or disabled .	All levels
Restart duration	Time period for complete reacquisition of OSPF neighbors.	All levels
Restart grace period	Time period for which the neighbors should consider the restarting routing device as part of the topology.	All levels

Table 117: show ospf overview Output Fields (*continued*)

Field name	Field Description	Level of Output
Graceful restart helper mode	(OSPFv2) Standard graceful restart helper capability (based on RFC 3623): enabled or disabled .	All levels
Restart-signaling helper mode	(OSPFv2) Restart signaling-based graceful restart helper capability (based on RFC 4811, RFC 4812, and RFC 4813): enabled or disabled .	All levels
Helper mode	(OSPFv3) Graceful restart helper capability: enabled or disabled .	All levels
Trace options	OSPF-specific trace options.	extensive
Trace file	Name of the file to receive the output of the tracing operation.	extensive
Area	Area number. Area 0.0.0.0 is the backbone area.	All levels
Stub type	Stub type of area: Normal Stub , Not Stub , or Not so Stubby Stub .	All levels
Authentication Type	Type of authentication: None , Password , or MD5 . NOTE: The Authentication Type field refers to the authentication configured at the <code>[edit protocols ospf area area-id]</code> level. Any authentication configured for an interface in this area will not affect the value of this field.	All levels
Area border routers	Number of area border routers.	All levels
Neighbors	Number of autonomous system boundary routers.	All levels

Sample Output

show ospf overview

```

user@host> show ospf overview
Instance: master
  Router ID: 10.255.245.6
  Route table index: 0
  Configured overload, expires in 118 seconds
  LSA refresh time: 50 minutes
  Restart: Enabled
    Restart duration: 20 sec
    Restart grace period: 40 sec
    Helper mode: enabled
  Area: 0.0.0.0
    Stub type: Not Stub
    Authentication Type: None
    Area border routers: 0, AS boundary routers: 0
    Neighbors
      Up (in full state): 0
  Topology: default (ID 0)
  Prefix export count: 0
  Full SPF runs: 1
  SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3

```

show ospf overview (With Database Protection)

```
user@host> show ospf overview
Instance: master
  Router ID: 10.255.112.218
  Route table index: 0
  LSA refresh time: 50 minutes
  Traffic engineering
  Restart: Enabled
    Restart duration: 180 sec
    Restart grace period: 210 sec
    Graceful restart helper mode: Enabled
    Restart-signaling helper mode: Enabled
  Database protection state: Normal
    Warning threshold: 70 percent
    Non self-generated LSAs: Current 582, Warning 700, Allowed 1000
    Ignore time: 30, Reset time: 60
    Ignore count: Current 0, Allowed 1
  Area: 0.0.0.0
    Stub type: Not Stub
    Authentication Type: None
    Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 160
  Topology: default (ID 0)
    Prefix export count: 0
    Full SPF runs: 70
    SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
    Backup SPF: Not Needed
```

show ospf3 overview (With Database Protection)

```
user@host> show ospf3 overview
Instance: master
  Router ID: 10.255.112.128
  Route table index: 0
  LSA refresh time: 50 minutes
  Database protection state: Normal
    Warning threshold: 80 percent
    Non self-generated LSAs: Current 3, Warning 8, Allowed 10
    Ignore time: 30, Reset time: 60
    Ignore count: Current 0, Allowed 2
  Area: 0.0.0.0
    Stub type: Not Stub
    Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 1
  Topology: default (ID 0)
    Prefix export count: 0
    Full SPF runs: 7
    SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
    Backup SPF: Not Needed
```

show ospf overview extensive

```
user@host> show ospf overview extensive
Instance: master
  Router ID: 1.1.1.103
  Route table index: 0
  Full SPF runs: 13, SPF delay: 0.200000 sec
  LSA refresh time: 50 minutes
```



```
Restart: Disabled
Trace options: lsa
Trace file: /var/log/ospf size 131072 files 10
Area: 0.0.0.0
  Stub type: Not Stub
  Authentication Type: None
  Area border routers: 0, AS boundary routers: 0
  Neighbors
    Up (in full state): 1
```

show (ospf | ospf3) route

Syntax	<pre>show (ospf ospf3) route <brief detail extensive> <abr asbr extern inter intra> <destination> <instance (default ipv4-multicast <i>instance-name</i>)> <logical-system (default ipv4-multicast <i>logical-system-name</i>)> <network> <no-backup-coverage> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)> <router> <topology (default ipv4-multicast <i>topology-name</i>)> <transit></pre>
Syntax (EX Series Switch and QFX Series)	<pre>show (ospf ospf3) route <brief detail extensive> <abr asbr extern inter intra> <destination> <instance <i>instance-name</i> <network> <no-backup-coverage> <router> <topology (default ipv4-multicast <i>topology-name</i>)> <transit></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>topology option introduced in Junos OS Release 9.0.</p> <p>realm option introduced in Junos OS Release 9.2.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Display the entries in the Open Shortest Path First (OSPF) routing table.
Options	<p>none—Display standard information about all entries in the OSPF routing table for all routing instances and all topologies.</p> <p>destination—Display routes to the specified IP address (with optional destination prefix length).</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>abr—(Optional) Display routes to area border routers.</p> <p>asbr—(Optional) Display routes to autonomous system border routers.</p> <p>extern—(Optional) Display external routes.</p> <p>inter—(Optional) Display interarea routes.</p> <p>intra—(Optional) Display intra-area routes.</p>

instance (default | ipv4-multicast | *instance-name*)—(Optional) Display entries for the default routing instance, the IPv4 multicast routing instance, or for the specified routing instance.

logical-system (default | ipv4-multicast | *logical-system-name*)—(Optional) Perform this operation on the default logical system, the IPv4 multicast logical system, or on a particular logical system.

network—(Optional) Display routes to networks.

no-backup-coverage—(Optional) Display routes with no backup coverage.

realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)—(OSPFv3 only) (Optional) Display entries in the routing table for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

router—(Optional) Display routes to all routers.

topology (default | ipv4-multicast | *topology-name*)—(OSPFv2 only) (Optional) Display routes for the default OSPF topology, IPv4 multicast topology, or for a particular topology.

transit—(Optional) (OSPFv3 only) Display OSPFv3 routes to pseudonodes.

Required Privilege Level

view

List of Sample Output

[show ospf route on page 755](#)
[show ospf route detail on page 755](#)
[show ospf3 route on page 755](#)
[show ospf3 route detail on page 756](#)
[show ospf route topology voice on page 756](#)

Output Fields

[Table 118 on page 753](#) list the output fields for the **show (ospf | ospf3) route** command. Output fields are listed in the approximate order in which they appear.

Table 118: show (ospf | ospf3) route Output Fields

Field Name	Field Description	Output Level
Topology	Name of the topology.	All levels
Prefix	Destination of the route.	All levels
Path type	How the route was learned: <ul style="list-style-type: none"> Inter—Interarea route Ext1—External type 1 route Ext2—External type 2 route Intra—Intra-area route 	All levels

Table 118: show (ospf | ospf3) route Output Fields (*continued*)

Field Name	Field Description	Output Level
Route type	The type of routing device from which the route was learned: <ul style="list-style-type: none"> • AS BR—Route to AS border router. • Area BR—Route to area border router. • Area/AS BR—Route to router that is both an Area BR and AS BR. • Network—Network router. • Router—Route to a router that is neither an Area BR nor an AS BR. • Transit—(OSPFv3 only) Route to a pseudonode representing a transit network, LAN, or nonbroadcast multiaccess (NBMA) link. • Discard—Route to a summary discard. 	All levels
NH Type	Next-hop type: LSP or IP .	All levels
Metric	Route's metric value.	All levels
NH-interface	(OSPFv3 only) Interface through which the route's next hop is reachable.	All levels
NH-addr	(OSPFv3 only) IPv6 address of the next hop.	All levels
NextHop Interface	(OSPFv2 only) Interface through which the route's next hop is reachable.	All levels
Nexthop addr/label	(OSPFv2 only) If the NH Type is IP , then it is the address of the next hop. If the NH Type is LSP , then it is the name of the label-switched path.	All levels
Area	Area ID of the route.	detail
Origin	Router from which the route was learned.	detail
Type 7	Route was learned through a not-so-stubby area (NSSA) link-state advertisement (LSA).	detail
P-bit	Route was learned through NSSA LSA and the propagate bit was set.	detail
Fwd NZ	Forwarding address is nonzero. Fwd NZ is only displayed if the route is learned through an NSSA LSA.	detail
optional-capability	Optional capabilities propagated in the router LSA. This field is in the output for intra-area router routes only (when Route Type is Area BR , AS BR , Area/AS BR , or Router), not for interarea router routes or network routes. Three bits in this field are defined as follows: <ul style="list-style-type: none"> • 0x4 (V)—Routing device is at the end of a virtual active link. • 0x2 (E)—Routing device is an autonomous system boundary router. • 0x1 (B)—Routing device is an area border router. 	detail

Table 118: show (ospf | ospf3) route Output Fields (*continued*)

Field Name	Field Description	Output Level
priority	The priority assigned to the prefix: <ul style="list-style-type: none"> • high • medium • low <p>NOTE: The priority field applies only to routes of type Network.</p>	detail

Sample Output

show ospf route

```

user@host> show ospf route
Prefix          Path   Route   NH   Metric  NextHop      Nexthop
                Type  Type    Type                Interface    addr/label
10.255.71.12    Intra Router  IP    1      fe-0/0/2.0  192.16.22.86
10.255.71.13/32 Intra Network IP    0      lo0.0
192.168.222.84/30 Intra Network LSP   1      fe-0/0/2.0  1sp-ab

```

show ospf route detail

```

user@host> show ospf route detail
Topology default Route Table:

Prefix          Path   Route   NH   Metric  NextHop      Nexthop
                Type  Type    Type                Interface    addr/label
10.255.14.174    Inter AS BR   IP    210    t1-3/0/1.0
  area 0.0.0.2, origin 10.255.14.185
10.255.14.178    Intra Router  IP    200    t3-3/1/3.0
  area 0.0.0.2, origin 10.255.14.178, optional-capability 0x0
10.210.1.0/30    Intra Network IP    10     t3-3/1/2.0
  area 0.0.0.2, origin 10.255.14.172, priority medium
100.1.1.1/32     Inter Network IP    210    t1-3/0/1.0
  area 0.0.0.2, origin 10.255.14.185, priority low
112.3.1.0/24     Ext2  Network  IP    0      t1-3/0/1.0
  area 0.0.0.0, origin 10.255.14.174, priority high
200.3.3.0/30     Inter Network IP    220    t1-3/0/1.0
  area 0.0.0.2, origin 10.255.14.185, priority high

```

show ospf3 route

```

user@host> show ospf3 route
Prefix          Path   Route   NH   Metric  NextHop      Nexthop
                Type  Type    Type                Interface    addr/label
10.255.71.13     Intra Router  IP    1      fe-0/0/2.0  192.16.22.86
  NH-interface fe-0/0/2.0, NH-addr fe80::290:69ff:fe9b:e002
10.255.71.13;0.0.0.2
10.255.245.1     Intra Router  IP    40     fxp1.1      192.168.36.17
  area 0.0.0.0, origin 10.255.245.1 optional-capability 0x0,
10.255.245.3     Intra AS BR   IP    1      fxp2.3      192.168.36.34
  area 0.0.0.0, origin 10.255.245.3 optional-capability 0x0,
10.255.245.1/32  Intra Network IP    40     fxp1.1      192.168.36.17

```

```

    area 0.0.0.0, origin 10.255.245.1, priority high
10.255.245.2/32      Intra Network   IP      0  lo0.0
    area 0.0.0.0, origin 10.255.245.2, priority medium
10.255.245.3/32      Intra Network   IP      1  fxp2.3      192.168.36.34

    area 0.0.0.0, origin 10.255.245.3, priority low
                        Intra Transit   IP      1
    NH-interface fe-0/0/2.0
192::168:222:84/126 Intra Network   IP      1
    NH-interface fe-0/0/2.0
abcd::71:12/128     Intra Network   IP      0
    NH-interface lo0.0
abcd::71:13/128     Intra Network   LSP     1
    NH-interface fe-0/0/2.0, NH-addr lsp-cd

```

show ospf3 route detail

```

user@host> show ospf3 route detail
Prefix                                Path   Route   NH   Metric
                                type  type   type
10.255.14.174                        Intra  Area/AS BR IP    110
    NH-interface so-1/2/2.0
    Area 0.0.0.0, Origin 10.255.14.174, Optional-capability 0x3
10.255.14.178                        Intra  Router  IP    200
    NH-interface t3-3/1/3.0
    Area 0.0.0.0, Origin 10.255.14.178, Optional-capability 0x0
10.255.14.185;0.0.0.2                Intra  Transit IP    200
    NH-interface t1-3/0/1.0
    NH-interface so-1/2/2.0
    Area 0.0.0.0, Origin 10.255.14.185
1000:1:1::1/128                      Inter  Network IP    110
    NH-interface so-1/2/2.0
    Area 0.0.0.0, Origin 10.255.14.174, Priority low
1001:2:1::/48                       Ext1   Network IP    110
    NH-interface so-1/2/2.0
    Area 0.0.0.0, Origin 10.255.14.174, Fwd NZ, Priority medium
1002:1:7::/48                       Ext2   Network IP    0
    NH-interface so-1/2/2.0
    Area 0.0.0.0, Origin 10.255.14.174, Fwd NZ, Priority low
1002:3:4::/48                       Ext2   Network IP    0
    NH-interface so-1/2/2.0
    Area 0.0.0.0, Origin 10.255.14.174, Fwd NZ, Priority high
abcd::10:255:14:172/128             Intra  Network IP    0
    NH-interface lo0.0
    Area 0.0.0.0, Origin 10.255.14.172, Priority low

```

show ospf route topology voice

```

user@host show ospf route topology voice
Topology voice Route Table:
Prefix      Path   Route   NH   Metric  NextHop      Nexthop
                                Type  Type   Type
10.255.8.2  Intra  Router  IP    1    so-0/2/0.0
10.255.8.3  Intra  Router  IP    2    so-0/2/0.0
10.255.8.1/32 Intra  Network IP    0    lo0.0
10.255.8.2/32 Intra  Network IP    1    so-0/2/0.0
10.255.8.3/32 Intra  Network IP    2    so-0/2/0.0
192.168.8.0/29 Intra  Network IP    2    so-0/2/0.0
192.168.8.44/30 Intra  Network IP    2    so-0/2/0.0
192.168.8.46/32 Intra  Network IP    1    so-0/2/0.0

```

192.168.8.48/30	Intra	Network	IP	1	so-0/2/1.0
192.168.8.52/30	Intra	Network	IP	2	so-0/2/0.0
192.168.9.44/30	Intra	Network	IP	1	so-0/2/0.0
192.168.9.45/32	Intra	Network	IP	2	so-0/2/0.0

show (ospf | ospf3) statistics

Syntax	show (ospf ospf3) statistics <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)> <realm (ipv4-multicast ipv4-unicast ipv6-multicast)>
Syntax (EX Series Switch and QFX Series)	show (ospf ospf3) statistics <instance <i>instance-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. realm option introduced in Junos OS Release 9.2. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	Display OSPF statistics.
Options	<p>none—Display OSPF statistics for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Display all statistics for the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>realm (ipv4-multicast ipv4-unicast ipv6-multicast)—(Optional) (OSPFv3 only) Display all statistics for the specified OSPFv3 realm, or address family. Use the realm option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> clear (ospf ospf3) statistics on page 693
List of Sample Output	show ospf statistics on page 760 show ospf statistics logical-system all on page 760 show ospf3 statistics on page 761
Output Fields	Table 119 on page 758 lists the output fields for the show (ospf ospf3) statistics command. Output fields are listed in the approximate order in which they appear.

Table 119: show (ospf | ospf3) statistics Output Fields

Field Name	Field Description
Packet type	Type of OSPF packet.
Total Sent/Total Received	Total number of packets sent and received.
Last 5 seconds Sent/Last 5 seconds Received	Total number of packets sent and received in the last 5 seconds.

Table 119: show (ospf | ospf3) statistics Output Fields (*continued*)

Field Name	Field Description
DBDs retransmitted	Total number of database description packets retransmitted, and number retransmitted in the last 5 seconds.
LSAs flooded	Total number of link-state advertisements flooded, and number flooded in the last 5 seconds.
LSAs flooded high-prio	<p>Total number of high priority link-state advertisements flooded, and number flooded in the last 5 seconds.</p> <p>A link-state advertisement is deemed a high priority if it has changed since it was last sent.</p>
LSAs retransmitted	Total number of link-state advertisements retransmitted, and number retransmitted in the last 5 seconds.
LSAs transmitted to nbr	Total number of link-state advertisements transmitted to a neighbor, and number transmitted in the last 5 seconds.
LSAs requested	Total number of link-state advertisements requested by neighboring devices, and number requested in the last 5 seconds.
LSAs acknowledged	Total number of link-state advertisements acknowledged, and number acknowledged in the last 5 seconds.
Flood queue depth	Total number of entries in the extended queue.
Total rexmit entries	Total number of retransmission entries waiting to be sent from the OSPF routing instance.
db summaries	Total number of database description summaries waiting to be sent from the OSPF routing instance.
lsreq entries	Total number of link-state request entries waiting to be sent from the OSPF routing instance.
Receive errors	<p>Number and type of receive errors. Some sample receive errors include:</p> <ul style="list-style-type: none"> • mtu mismatches • no interface found • no virtual link found • nssa mismatches • stub area mismatches • subnet mismatches <p>If there are no receive errors, the output displays none.</p>

Sample Output

show ospf statistics

```

user@host> show ospf statistics
Packet type          Total
                   Sent      Received
Hello                31        14
DbD                  9         10
LSReq                2          2
LSUpdate             8         16
LSAck                9          9
                   Last 5 seconds
                   Sent      Received
Hello                2          2
DbD                  0          0
LSReq                0          0
LSUpdate             0          0
LSAck                0          0

DBDs retransmitted   :          3, last 5 seconds :          0
LSAs flooded         :         12, last 5 seconds :          0
LSAs flooded high-prio :          0, last 5 seconds :          0
LSAs retransmitted   :          0, last 5 seconds :          0
LSAs transmitted to nbr:          3, last 5 seconds :          0
LSAs requested       :          5, last 5 seconds :          0
LSAs acknowledged    :         19, last 5 seconds :          0

Flood queue depth    :          0
Total rexmit entries :          0
db summaries         :          0
lsreq entries        :          0

Receive errors:
  862 no interface found
 115923 no virtual link found

```

show ospf statistics logical-system all

```

user@host> show ospf statistics logical-system all
logical-system: C
OSPF instance is not running
-----

logical-system: B
Packet type          Total
                   Sent      Received
Hello              313740      313653
DbD                 3          2
LSReq              1          1
LSUpdate           2752      1825
LSAck              1821      2747
                   Last 5 seconds
                   Sent      Received
Hello               1          0
DbD                 0          0
LSReq               0          0
LSUpdate            0          0
LSAck               0          0

DBDs retransmitted   :          0, last 5 seconds :          0
LSAs flooded         :        2741, last 5 seconds :          0
LSAs flooded high-prio :         10, last 5 seconds :          0
LSAs retransmitted   :          0, last 5 seconds :          0
LSAs transmitted to nbr:          2, last 5 seconds :          0
LSAs requested       :          1, last 5 seconds :          0
LSAs acknowledged    :       1831, last 5 seconds :          0

Flood queue depth    :          0
Total rexmit entries :          0
db summaries         :          0
lsreq entries        :          0

Receive errors:

```

```

None
-----

logical-system: A

Packet type          Total          Last 5 seconds
                   Sent      Received      Sent      Received
Hello                313698      313695         0         0
  DbD                  2         3         0         0
  LSReq                1         1         0         0
LSUpdate             1825      2752         0         0
LSAck                2747      1821         0         0

DBDs retransmitted   :          0, last 5 seconds :          0
LSAs flooded         :        1825, last 5 seconds :          0
LSAs flooded high-prio :         10, last 5 seconds :          0
LSAs retransmitted   :          0, last 5 seconds :          0
LSAs transmitted to nbr:         1, last 5 seconds :          0
LSAs requested       :          2, last 5 seconds :          0
LSAs acknowledged   :        2748, last 5 seconds :          0

Flood queue depth    :          0
Total rexmit entries :          0
db summaries         :          0
lsreq entries        :          0

Receive errors:
None
-----

```

show ospf3 statistics

```

user@host> show ospf3 statistics

Packet type          Total          Last 5 seconds
                   Sent      Received      Sent      Received
Hello                0         0         0         0
  DbD                  0         0         0         0
  LSReq                0         0         0         0
LSUpdate             0         0         0         0
LSAck                0         0         0         0

DBDs retransmitted   :          0, last 5 seconds :          0
LSAs flooded         :          0, last 5 seconds :          0
LSAs flooded high-prio :          0, last 5 seconds :          0
LSAs retransmitted   :          0, last 5 seconds :          0
LSAs transmitted to nbr:          0, last 5 seconds :          0
LSAs requested       :          0, last 5 seconds :          0
LSAs acknowledged   :          0, last 5 seconds :          0

Flood queue depth    :          0
Total rexmit entries :          0
db summaries         :          0
lsreq entries        :          0

Receive errors:
None

```


CHAPTER 16

Protocol-Independent Routing Operational Mode Commands

- `show as-path`
- `show as-path domain`
- `show as-path summary`
- `show route`
- `show route active-path`
- `show route advertising-protocol`
- `show route all`
- `show route aspath-regex`
- `show route best`
- `show route brief`
- `show route ccc`
- `show route community`
- `show route community-name`
- `show route damping`
- `show route detail`
- `show route exact`
- `show route export`
- `show route export vrf-target`
- `show route extensive`
- `show route flow validation`
- `show route forwarding-table`
- `show route forwarding-table interface-name`
- `show route hidden`
- `show route inactive-path`
- `show route inactive-prefix`
- `show route instance`

- `show route label`
- `show route label-switched-path`
- `show route localization`
- `show route martians`
- `show route match-prefix`
- `show route next-hop`
- `show route no-community`
- `show route output`
- `show route protocol`
- `show route range`
- `show route receive-protocol`
- `show route resolution`
- `show route snooping`
- `show route source-gateway`
- `show route summary`
- `show route table`
- `show route terse`

show as-path

Syntax	show as-path <brief detail> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show as-path <brief detail>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	<p>Display the distribution of autonomous system (AS) paths that the local routing device is using (usually through the routing table). Use this command to debug problems for AS paths and to understand how AS paths have been manipulated through a policy (through the as-path-prepend action) or through aggregation.</p> <p>AS paths are stored in a hash table. A hash table is one method for fast lookup. Each entry in the table is called a bucket. Junos OS computes a hash value that indicates in which bucket the AS path is stored. The AS paths are dispersed among the hash buckets so that a manageable number of AS paths is stored in each bucket. Only unique AS paths are stored. Duplicate AS paths increase a reference count, but do not increase the number of AS paths stored in the hash table.</p>
Options	<p>none—Display basic information about AS paths that the local routing device is using (same as brief).</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show as-path summary on page 771
List of Sample Output	show as-path on page 766 show as-path detail on page 767
Output Fields	Table 120 on page 765 lists the output fields for the show as-path command. Output fields are listed in the approximate order in which they appear.

Table 120: show as-path Output Fields

Field Name	Field Description	Level of Output
Total AS paths	Total number of AS paths.	brief none
Bucket	Bucket number.	All levels

Table 120: show as-path Output Fields (*continued*)

Field Name	Field Description	Level of Output
Count	Number of AS path entries in this bucket.	All levels
AS path	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • ?—Incomplete; typically, the AS path was aggregated. • Atomic—Route is an aggregate of several route prefixes. • Aggregator—Routing device has summarized a range of prefixes. 	All levels
domain	Number of independent AS domains. The AS paths of an independent AS domain are not shared with the AS paths and AS path attributes of other domains, including the master routing instance domain.	detail
neighbor as	AS peer address.	detail
length	Length of the AS path.	detail
segments	Length of the AS segment descriptor.	detail
references	Path reference count.	detail

Sample Output

show as-path

```

user@host> show as-path
Total AS paths: 30382
Bucket 0      Count: 36
  I
  14203 2914 174 31752 I
  14203 2914 701 21512 I
  14203 2914 1239 26632 I
  14203 2914 1239 29704 I
  14203 2914 4323 10248 I
  14203 2914 4766 23560 I
  14203 2914 6395 32776 I
  14203 2914 7911 11272 I
  14203 2914 12180 18440 I
  14203 2914 17408 17416 I
  14203 2914 701 702 24586 I
  14203 2914 1239 4657 9226 I
  14203 2914 1239 7132 16394 I
  14203 2914 1299 8308 34826 I
  14203 2914 3320 5603 28682 I
  14203 2914 3491 1680 33802 I
  14203 2914 3549 7908 27658 I
  14203 2914 3549 20804 30730 I
  14203 2914 7018 2687 9226 I
  14203 2914 174 9318 9318 23564 I

```



```

14203 2914 701 3786 3786 23564 I
14203 2914 701 4761 4795 9228 I
14203 2914 1239 7132 5673 18444 I
14203 2914 3491 20485 24588 24588 I
14203 2914 5511 2200 1945 2060 I
14203 2914 7911 14325 14325 14348 I
14203 2914 701 4637 9230 9230 9230 I
14203 2914 6395 14 14 14 14 I
14203 2914 9299 6163 6163 6163 6163 9232 I
14203 2914 3356 3356 3356 3356 3356 11955 21522 I
14203 2914 9837 9837 9219 I Aggregator: 9219 202.27.91.253
14203 2914 174 30209 30222 30222 30222 ?
14203 2914 1299 5377 I (Atomic) Aggregator: 5377 193.219.192.22
14203 2914 4323 36097 I (Atomic) Aggregator: 36097 216.69.252.254
14203 2914 209 2516 17676 23813 I (Atomic) Aggregator: 23813 219.127.233.66
Bucket 1    Count: 28
14203 2914 35847 I
14203 2914 174 19465 I
14203 2914 174 35849 I
14203 2914 2828 32777 I
14203 2914 4323 14345 I
14203 2914 4323 29705 I
14203 2914 6395 32777 I
...

```

show as-path detail

```

user@host> show as-path detail
Total AS paths: 30410
Bucket 0    Count: 36
AS path: I
  domain 0, length 0, segments 0, references 54
AS path: 14203 2914 174 31752 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 701 21512 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 1239 26632 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 1239 29704 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 4323 10248 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 4766 23560 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 6395 32776 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 3
AS path: 14203 2914 7911 11272 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 2
AS path: 14203 2914 12180 18440 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 3
AS path: 14203 2914 17408 17416 I
  domain 1, neighbor as: 14203, length 4, segments 1, references 3
AS path: 14203 2914 701 702 24586 I
  domain 1, neighbor as: 14203, length 5, segments 1, references 3
AS path: 14203 2914 1239 4657 9226 I
  domain 1, neighbor as: 14203, length 5, segments 1, references 7
AS path: 14203 2914 1239 7132 16394 I
  domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 1299 8308 34826 I
  domain 1, neighbor as: 14203, length 5, segments 1, references 2

```

```
AS path: 14203 2914 3320 5603 28682 I
  domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 3491 1680 33802 I
  domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 3549 7908 27658 I
  domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 3549 20804 30730 I
  domain 1, neighbor as: 14203, length 5, segments 1, references 2
AS path: 14203 2914 7018 2687 9226 I
  domain 1, neighbor as: 14203, length 5, segments 1, references 3
AS path: 14203 2914 174 9318 9318 23564 I
  domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 701 3786 3786 23564 I
  domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 701 4761 4795 9228 I
  domain 1, neighbor as: 14203, length 6, segments 1, references 14
AS path: 14203 2914 1239 7132 5673 18444 I
  domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 3491 20485 24588 24588 I
  domain 1, neighbor as: 14203, length 6, segments 1, references 4
AS path: 14203 2914 5511 2200 1945 2060 I
  domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 7911 14325 14325 14348 I
  domain 1, neighbor as: 14203, length 6, segments 1, references 2
AS path: 14203 2914 701 4637 9230 9230 9230 I
  domain 1, neighbor as: 14203, length 7, segments 1, references 3
AS path: 14203 2914 6395 14 14 14 14 I
  domain 1, neighbor as: 14203, length 7, segments 1, references 10
...
```

show as-path domain

Syntax	show as-path domain <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show as-path domain
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display autonomous system (AS) path domain information.
Options	none —(Optional) Display AS path domain information for all routing instances. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show as-path domain on page 770
Output Fields	Table 121 on page 769 lists the output fields for the show as-path domain command. Output fields are listed in the approximate order in which they appear

Table 121: show as-path domain Output Fields

Field Name	Field Description
Domain	Number of independent AS domains. The AS paths of an independent AS domain are not shared with the AS paths and AS path attributes of other domains, including the master routing instance domain.
Primary	Primary AS number.
References	Path reference count.
Number Paths	Number of known AS paths.
Flags	Information about the AS path: <ul style="list-style-type: none"> • ASLoop—Path contains an AS loop. • Atomic—Path includes the ATOMIC_AGGREGATE path attribute. • Local—Path was created by local aggregation. • Master—Path was created by the master routing instance.
Local AS	AS number of the local routing device.
Loops	How many times this AS number can appear in an AS path.

Sample Output

show as-path domain

```
user@host> show as-path domain
Domain: 1          Primary: 10458
References:        3 Paths:      30383
Flags: Master
Local AS: 10458   Loops: 1
```

show as-path summary

Syntax	show as-path summary <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show as-path summary
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display autonomous system (AS) path summary information. AS paths are stored in a hash table. A hash table is one method for fast lookup. Each entry in the table is called a bucket. Junos OS computes a hash value that indicates in which bucket the AS path is stored. The AS paths are dispersed among the hash buckets so that a manageable number of AS paths is stored in each bucket. Only unique AS paths are stored. Duplicate AS paths increase a reference count, but do not increase the number of AS paths stored in the hash table.
Options	none —(Optional) Display AS path summary information for all routing instances. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show as-path on page 765
List of Sample Output	show as-path summary on page 772
Output Fields	Table 122 on page 771 lists the output fields for the show as-path summary command. Output fields are listed in the approximate order in which they appear.

Table 122: show as-path summary Output Fields

Field Name	Field Description
AS Paths	Number of AS paths.
Buckets	Number of hash buckets in use.
Max	Maximum number of AS path entries per bucket.
Min	Minimum number of AS path entries per bucket.
Avg	Average number of AS path entries per bucket.
Std deviation	Standard deviation of AS path entries per bucket.

Sample Output

show as-path summary

```
user@host> show as-path summary
AS Paths Buckets Max Min Avg Std deviation
30425    1024    95  12  29  6.481419
```

show route

Syntax	<pre>show route <all> <destination-prefix> <logical-system (all logical-system-name)> <private></pre>
Syntax (EX Series Switches)	<pre>show route <all> <destination-prefix> <private></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Option private introduced in Junos OS Release 9.5.</p> <p>Option private introduced in Junos OS Release 9.5 for EX Series switches.</p>
Description	Display the active entries in the routing tables.
Options	<p>none—Display brief information about all active entries in the routing tables.</p> <p>all—(Optional) Display information about all routing tables, including private, or internal, routing tables.</p> <p>destination-prefix—(Optional) Display active entries for the specified address or range of addresses.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>private—(Optional) Display information only about all private, or internal, routing tables.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring RIP</i> • <i>Example: Configuring RIPng</i> • <i>Example: Configuring IS-IS</i> • <i>Examples: Configuring Internal BGP Peering</i> • <i>Examples: Configuring External BGP Peering</i> • <i>Examples: Configuring OSPF Routing Policy</i>
List of Sample Output	<p>show route on page 776</p> <p>show route on page 776</p> <p>show route destination-prefix on page 777</p> <p>show route extensive on page 777</p>

Output Fields Table 123 on page 774 describes the output fields for the **show route** command. Output fields are listed in the approximate order in which they appear.

Table 123: show route Output Fields

Field Name	Field Description
<i>routing-table-name</i>	Name of the routing table (for example, inet.0).
<i>number destinations</i>	Number of destinations for which there are routes in the routing table.
<i>number routes</i>	<p>Number of routes in the routing table and total number of routes in the following states:</p> <ul style="list-style-type: none"> • active (routes that are active). • holddown (routes that are in the pending state before being declared inactive). A holddown route was once the active route and is no longer the active route. The route is in the holddown state because a protocol still has interest in the route, meaning that the interest bit is set. A protocol might have its interest bit set on the previously active route because the protocol is still advertising the route. The route will be deleted after all protocols withdraw their advertisement of the route and remove their interest bit. A persistent holddown state often means that the interested protocol is not releasing its interest bit properly. <p>However, if you have configured advertisement of multiple routes (with the add-path or advertise-inactive statement), the holddown bit is most likely set because BGP is advertising the route as an active route. In this case, you can ignore the holddown state because nothing is wrong.</p> <ul style="list-style-type: none"> • hidden (routes that are not used because of a routing policy).
<i>destination-prefix</i>	<p>Route destination (for example:10.0.0.1/24). Sometimes the route information is presented in another format, such as:</p> <ul style="list-style-type: none"> • MPLS-label (for example, 80001). • interface-name (for example, ge-1/0/2). • neighbor-address:control-word-status:encapsulation type:vc-id:source (Layer 2 circuit only. For example, 10.1.1.195:NoCtrlWord:1:1:Local/96): <ul style="list-style-type: none"> • neighbor-address—Address of the neighbor. • control-word-status—Whether the use of the control word has been negotiated for this virtual circuit: NoCtrlWord or CtrlWord. • encapsulation type—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport. • vc-id—Virtual circuit identifier. • source—Source of the advertisement: Local or Remote.
[protocol, preference]	<p>Protocol from which the route was learned and the preference value for the route.</p> <ul style="list-style-type: none"> • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • -—A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route. <p>In every routing metric except for the BGP LocalPref attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the LocalPref value in the Preference2 field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the LocalPref value for Route 2 is 155, the Preference2 value is -156. Route 2 is preferred because it has a higher LocalPref value and a lower Preference2 value.</p>

Table 123: show route Output Fields (*continued*)

Field Name	Field Description
<i>weeks:days</i> <i>hours:minutes:seconds</i>	How long the route been known (for example, 2w4d 13:11:14 , or 2 weeks, 4 days, 13 hours, 11 minutes, and 14 seconds).
metric	Cost value of the indicated route. For routes within an AS, the cost is determined by the IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
localpref	Local preference value included in the route.
from	Interface from which the route was received.
AS path	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • ?—Incomplete; typically, the AS path was aggregated. <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> • []—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured. • { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • ()—Parentheses enclose a confederation. • ([])—Parentheses and brackets enclose a confederation set. <p>NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p>
validation-state	<p>(BGP-learned routes) Validation status of the route:</p> <ul style="list-style-type: none"> • Invalid—Indicates that the prefix is found, but either the corresponding AS received from the EBGp peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • Unknown—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • Unverified—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers. • Valid—Indicates that the prefix and autonomous system pair are found in the database.
to	<p>Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.</p> <p>If the destination is Discard, traffic is dropped.</p>

Table 123: show route Output Fields (*continued*)

Field Name	Field Description
via	<p>Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word Selected. This field can also contain the following information:</p> <ul style="list-style-type: none"> • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing. • lsp-path-name—Name of the LSP used to reach the next hop. • label-action—MPLS label and operation occurring at the next hop. The operation can be pop (where a label is removed from the top of the stack), push (where another label is added to the label stack), or swap (where a label is replaced by another label). For VPNs, expect to see multiple push operations, corresponding to the inner and outer labels required for VPN routes (in the case of a direct PE-to-PE connection, the VPN route would have the inner label push only).

Sample Output

show route

```

user@host> show route
inet.0: 11 destinations, 12 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:65500:1:10.0.0.20/240
    *[MVPN/70] 19:53:41, metric2 1
    Indirect
1:65500:1:10.0.0.40/240
    *[BGP/170] 19:53:29, localpref 100, from 10.0.0.30
    AS path: I
    > to 10.0.24.4 via lt-0/3/0.24, label-switched-path toD
    [BGP/170] 19:53:26, localpref 100, from 10.0.0.33
    AS path: I
    > to 10.0.24.4 via lt-0/3/0.24, label-switched-path toD
1:65500:1:10.0.0.60/240
    *[BGP/170] 19:53:29, localpref 100, from 10.0.0.30
    AS path: I
    > to 10.0.28.8 via lt-0/3/0.28, label-switched-path toF
    [BGP/170] 19:53:25, localpref 100, from 10.0.0.33
    AS path: I
    > to 10.0.28.8 via lt-0/3/0.28, label-switched-path toF

```

show route

The following sample output shows a VPN route with composite next hops enabled. The first **Push** operation corresponds to the outer label. The second **Push** operation corresponds to the inner label.

```

user@host> show route 70.0.0.0

13979:665001.inet.0: 871 destinations, 3556 routes (871 active, 0 holddown, 0
hidden)

```

+ = Active Route, - = Last Active, * = Both

```
70.0.0.0/24      @[BGP/170] 00:28:32, localpref 100, from 10.9.9.160
                  AS path: 13980 ?, validation-state: unverified
                  > to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)
                  [BGP/170] 00:28:28, localpref 100, from 10.9.9.169
                  AS path: 13980 ?, validation-state: unverified
                  > to 10.100.0.42 via ae2.0, Push 126016, Push 300368(top)
                  #[Multipath/255] 00:28:28, metric2 102
                  > to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)
                  to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)
```

show route destination-prefix

```
user@host> show route 172.16.0.0/12
```

```
inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
172.16.0.0/12      *[Static/5] 2w4d 12:54:27
                  > to 192.168.167.254 via fxp0.0
```

show route extensive

```
user@host> show route extensive
```

```
v1.mvpn.0: 5 destinations, 8 routes (5 active, 1 holddown, 0 hidden)
1:65500:1:10.0.0.40/240 (1 entry, 1 announced)
  *BGP Preference: 170/-101
    PMSI: Flags 0x0: Label[0:0:0]: PIM-SM: Sender 10.0.0.40 Group 225.1.1.1

    Next hop type: Indirect
    Address: 0x92455b8
    Next-hop reference count: 2
    Source: 10.0.0.30
    Protocol next hop: 10.0.0.40
    Indirect next hop: 2 no-forward
    State: <Active Int Ext>
      Local AS: 65500 Peer AS: 65500
    Age: 3 Metric2: 1
    Validation State: unverified
    Task: BGP_65500.10.0.0.30+179
    Announcement bits (2): 0-PIM.v1 1-mvpn global task
    AS path: I (Originator) Cluster list: 10.0.0.30
    AS path: Originator ID: 10.0.0.40
    Communities: target:65520:100
    Import Accepted
    Localpref: 100
    Router ID: 10.0.0.30
    Primary Routing Table bgp.mvpn.0
    Indirect next hops: 1
      Protocol next hop: 10.0.0.40 Metric: 1
      Indirect next hop: 2 no-forward
      Indirect path forwarding next hops: 1
        Next hop type: Router
        Next hop: 10.0.24.4 via lt-0/3/0.24 weight 0x1
      10.0.0.40/32 Originating RIB: inet.3
        Metric: 1 Node path count: 1
        Forwarding nexthops: 1
          Nexthop: 10.0.24.4 via lt-0/3/0.24
```

show route active-path

Syntax	show route active-path <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show route active-path <brief detail extensive terse>
Release Information	Command introduced in Junos OS Release 8.0. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display all active routes for destinations. An active route is a route that is selected as the best path. Inactive routes are not displayed.
Options	none —Display all active routes. brief detail extensive terse —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief . logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route active-path on page 778 show route active-path brief on page 779 show route active-path detail on page 779 show route active-path extensive on page 780 show route active-path terse on page 782
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route active-path

```
user@host> show route active-path

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.70.19/32    *[Direct/0] 21:33:52
                  > via lo0.0
10.255.71.50/32    *[IS-IS/15] 00:18:13, metric 10
                  > to 100.1.2.1 via so-2/1/3.0
100.1.2.0/24      *[Direct/0] 00:18:36
                  > via so-2/1/3.0
100.1.2.2/32      *[Local/0] 00:18:41
                  Local via so-2/1/3.0
192.168.64.0/21   *[Direct/0] 21:33:52
```

```

> via fxp0.0
192.168.70.19/32  *Local/0] 21:33:52
                  Local via fxp0.0

```

show route active-path brief

The output for the **show route active-path brief** command is identical to that for the **show route active-path** command. For sample output, see [show route active-path on page 778](#).

show route active-path detail

```

user@host> show route active-path detail

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)

10.255.70.19/32 (1 entry, 1 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 3
    Next hop: via lo0.0, selected
    State: <Active Int>
    Local AS: 200
    Age: 21:37:10
    Task: IF
    Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

    AS path: I

10.255.71.50/32 (1 entry, 1 announced)
  *IS-IS Preference: 15
    Level: 1
    Next hop type: Router, Next hop index: 397
    Next-hop reference count: 4
    Next hop: 100.1.2.1 via so-2/1/3.0, selected
    State: <Active Int>
    Local AS: 200
    Age: 21:31 Metric: 10
    Task: IS-IS
    Announcement bits (4): 0-KRT 2-IS-IS 5-Resolve tree 2 6-Resolve
tree 3
    AS path: I

100.1.2.0/24 (1 entry, 1 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 3
    Next hop: via so-2/1/3.0, selected
    State: <Active Int>
    Local AS: 200
    Age: 21:54
    Task: IF
    Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

    AS path: I

100.1.2.2/32 (1 entry, 1 announced)
  *Local Preference: 0
    Next hop type: Local
    Next-hop reference count: 11
    Interface: so-2/1/3.0
    State: <Active NoReadvrt Int>

```

```
Local AS: 200
Age: 21:59
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

192.168.64.0/21 (1 entry, 1 announced)
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 3
Next hop: via fxp0.0, selected
State: <Active Int>
Local AS: 200
Age: 21:37:10
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

192.168.70.19/32 (1 entry, 1 announced)
*Local Preference: 0
Next hop type: Local
Next-hop reference count: 11
Interface: fxp0.0
State: <Active NoReadvrt Int>
Local AS: 200
Age: 21:37:10
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I
```

show route active-path extensive

```
user@host> show route active-path extensive

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
10.255.70.19/32 (1 entry, 1 announced)
TSI:
IS-IS level 1, LSP fragment 0
IS-IS level 2, LSP fragment 0
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 3
Next hop: via lo0.0, selected
State: <Active Int>
Local AS: 200
Age: 21:39:47
Task: IF
Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3
AS path: I

10.255.71.50/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.255.71.50/32 -> {100.1.2.1}
IS-IS level 2, LSP fragment 0
*IS-IS Preference: 15
Level: 1
Next hop type: Router, Next hop index: 397
Next-hop reference count: 4
Next hop: 100.1.2.1 via so-2/1/3.0, selected
State: <Active Int>
```

```

Local AS: 200
Age: 24:08 Metric: 10
Task: IS-IS
Announcement bits (4): 0-KRT 2-IS-IS 5-Resolve tree 2 6-Resolve
tree 3
AS path: I

100.1.2.0/24 (1 entry, 1 announced)
TSI:
IS-IS level 1, LSP fragment 0
IS-IS level 2, LSP fragment 0
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 3
Next hop: via so-2/1/3.0, selected
State: <Active Int>
Local AS: 200
Age: 24:31
Task: IF
Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3
AS path: I

100.1.2.2/32 (1 entry, 1 announced)
*Local Preference: 0
Next hop type: Local
Next-hop reference count: 11
Interface: so-2/1/3.0
State: <Active NoReadvrt Int>
Local AS: 200
Age: 24:36
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

192.168.64.0/21 (1 entry, 1 announced)
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 3
Next hop: via fxp0.0, selected
State: <Active Int>
Local AS: 200
Age: 21:39:47
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

192.168.70.19/32 (1 entry, 1 announced)
*Local Preference: 0
Next hop type: Local
Next-hop reference count: 11
Interface: fxp0.0
State: <Active NoReadvrt Int>
Local AS: 200
Age: 21:39:47
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I

```

show route active-path terse

```
user@host> show route active-path terse
```

```
inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
```

```
+ = Active Route, - = Last Active, * = Both
```

A	Destination	P	Prf	Metric 1	Metric 2	Next hop	AS path
*	10.255.70.19/32	D	0			>1o0.0	
*	10.255.71.50/32	I	15	10		>100.1.2.1	
*	100.1.2.0/24	D	0			>so-2/1/3.0	
*	100.1.2.2/32	L	0			Local	
*	192.168.64.0/21	D	0			>fxp0.0	
*	192.168.70.19/32	L	0			Local	

show route advertising-protocol

Syntax	<code>show route advertising-protocol <i>protocol neighbor-address</i></code> <code><brief detail extensive terse></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display the routing information as it has been prepared for advertisement to a particular neighbor of a particular dynamic routing protocol.
Options	<p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>neighbor-address</i>—Address of the neighboring router to which the route entry is being transmitted.</p> <p><i>protocol</i>—Protocol transmitting the route:</p> <ul style="list-style-type: none"> • bgp—Border Gateway Protocol • dvmrp—Distance Vector Multicast Routing Protocol • msdp—Multicast Source Discovery Protocol • pim—Protocol Independent Multicast • rip—Routing Information Protocol • ripng—Routing Information Protocol next generation
Additional Information	Routes displayed are routes that the routing table has exported into the routing protocol and that have been filtered by the associated protocol's export routing policy statements.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring the MED Attribute Directly</i>
List of Sample Output	show route advertising-protocol bgp (Layer 3 VPN) on page 785 show route advertising-protocol bgp detail on page 786 show route advertising-protocol bgp detail (Layer 2 VPN) on page 786 show route advertising-protocol bgp detail (Layer 3 VPN) on page 786 show route advertising-protocol bgp extensive all (Next Hop Self with RIB-out IP Address) on page 786
Output Fields	Table 124 on page 784 lists the output fields for the show route advertising-protocol command. Output fields are listed in the approximate order in which they appear.

Table 124: show route advertising-protocol Output Fields

Field Name	Field Description	Level of Output
<i>routing-table-name</i>	Name of the routing table—for example, inet.0.	All levels
<i>number destinations</i>	Number of destinations for which there are routes in the routing table.	All levels
<i>number routes</i>	Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> • active (routes that are active) • holddown (routes that are in the pending state before being declared inactive) • hidden (routes that are not used because of a routing policy) 	All levels
Prefix	Destination prefix.	brief none
<i>destination-prefix (entry, announced)</i>	Destination prefix. The entry value is the number of routes for this destination, and the announced value is the number of routes being announced for this destination.	detail extensive
BGP group and type	BGP group name and type (Internal or External).	detail extensive
Route Distinguisher	Unique 64-bit prefix augmenting each IP subnet.	detail extensive
Advertised Label	Incoming label advertised by the LDP. When an IP packet enters a label-switched path (LSP), the ingress router examines the packet and assigns it a label based on its destination, placing the label in the packet's header. The label transforms the packet from one that is forwarded based on its IP routing information to one that is forwarded based on information associated with the label.	detail extensive
Label-Base, range	First label in a block of labels and label block size. A remote PE router uses this first label when sending traffic toward the advertising PE router.	detail extensive
VPN Label	Virtual private network (VPN) label. Packets are sent between CE and PE routers by advertising VPN labels. VPN labels transit over either an RSVP or an LDP LSP tunnel.	detail extensive
Nexthop	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route. If the next-hop advertisement to the peer is Self , and the RIB-out next hop is a specific IP address, the RIB-out IP address is included in the extensive output. See show route advertising-protocol bgp extensive all (Next Hop Self with RIB-out IP Address) on page 786.	All levels
MED	Multiple exit discriminator value included in the route.	brief
Lclpref or Localpref	Local preference value included in the route.	All levels

Table 124: show route advertising-protocol Output Fields (*continued*)

Field Name	Field Description	Level of Output
AS path	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • ?—Incomplete; typically, the AS path was aggregated. <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> • []—Brackets enclose the local AS number associated with the AS path if configured on the router, or if AS path prepending is configured. • { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • ()—Parentheses enclose a confederation. • ([])—Parentheses and brackets enclose a confederation set. <p>NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p>	All levels
Communities	Community path attribute for the route. See the output field table for the show route detail command for all possible values for this field.	detail extensive
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.	detail extensive
Attrset AS	Number, local preference, and path of the autonomous system (AS) that originated the route. These values are stored in the Attrset attribute at the originating router.	detail extensive
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).	detail extensive
control flags	Control flags: none or Site Down .	detail extensive
mtu	Maximum transmission unit (MTU) of the Layer 2 circuit.	detail extensive

Sample Output

show route advertising-protocol bgp (Layer 3 VPN)

```

user@host> show route advertising-protocol bgp 10.255.14.171
VPN-A.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
10.255.14.172/32 Self              1      100 I
VPN-B.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
10.255.14.181/32 Self              2      100 I

```

show route advertising-protocol bgp detail

```
user@host> show route advertising-protocol bgp 111.222.1.3 detail
bgp20.inet.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
111.222.1.11/32 (1 entry, 1 announced)
  BGP group pe-pe type Internal
    Route Distinguisher: 111.255.14.11:69
    Advertised Label: 100000
    next hop: Self
    Localpref: 100
    AS path: 2 I
    Communities: target:69:20
    AIGP 210
111.8.0.0/16 (1 entry, 1 announced)
  BGP group pe-pe type Internal
    Route Distinguisher: 111.255.14.11:69
    Advertised Label: 100000
    Next hop: Self
    Localpref: 100
    AS path: 2 I
    Communities: target:69:20
    AIGP 210
```

show route advertising-protocol bgp detail (Layer 2 VPN)

```
user@host> show route advertising-protocol bgp 192.168.24.1 detail
vpn-a.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
192.168.16.1:1:1:1/96 (1 entry, 1 announced)
  BGP group int type Internal
    Route Distinguisher: 192.168.16.1:1
    Label-base : 32768, range : 3
    Nexthop: Self
    Localpref: 100
    AS path: I
    Communities: target:65412:100
    AIGP 210
    Layer2-info: encaps:VLAN, control flags:, mtu:
```

show route advertising-protocol bgp detail (Layer 3 VPN)

```
user@host> show route advertising-protocol bgp 10.255.14.176 detail
vpna.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
* 10.49.0.0/30 (1 entry, 1 announced)
  BGP group ibgp type Internal
    Route Distinguisher: 10.255.14.174:2
    VPN Label: 101264
    Nexthop: Self
    Localpref: 100
    AS path: I
    Communities: target:200:100
    AIGP 210
    AttrSet AS: 100
      Localpref: 100
      AS path: I
  ...
```

show route advertising-protocol bgp extensive all (Next Hop Self with RIB-out IP Address)

```
user@host> show route advertising-protocol bgp 200.0.0.2 170.0.1.0/24 extensive all
inet.0: 13 destinations, 19 routes (13 active, 0 holddown, 6 hidden)
  170.0.1.0/24 (2 entries, 1 announced)
```

```

BGP group eBGP-INTEROP type External
  Nexthop: Self (rib-out 10.100.3.2)
  AS path: [4713] 200 I
...
```

show route all

Syntax	<code>show route all</code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches)	<code>show route all</code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display information about all routes in all routing tables, including private, or internal, tables.
Options	none —Display information about all routes in all routing tables, including private, or internal, tables. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route all on page 788
Output Fields	In Junos OS Release 9.5 and later, only the output fields for the show route all command display all routing tables, including private, or hidden, routing tables. The output field table of the show route command does not display entries for private, or hidden, routing tables in Junos OS Release 9.5 and later.

Sample Output

show route all

The following example displays a snippet of output from the **show route** command and then displays the same snippet of output from the **show route all** command:

```
user@host> show route
mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
0          *[MPLS/0] 2d 02:24:39, metric 1
            Receive
1          *[MPLS/0] 2d 02:24:39, metric 1
            Receive
2          *[MPLS/0] 2d 02:24:39, metric 1
            Receive
800017     *[VPLS/7] 1d 14:00:16
            > via vt-3/2/0.32769, Pop
800018     *[VPLS/7] 1d 14:00:26
            > via vt-3/2/0.32772, Pop

user@host> show route all
mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
```

```
+ = Active Route, - = Last Active, * = Both
0          *[MPLS/0] 2d 02:19:12, metric 1
           Receive
1          *[MPLS/0] 2d 02:19:12, metric 1
           Receive
2          *[MPLS/0] 2d 02:19:12, metric 1
           Receive
800017     *[VPLS/7] 1d 13:54:49
           > via vt-3/2/0.32769, Pop
800018     *[VPLS/7] 1d 13:54:59
           > via vt-3/2/0.32772, Pop
vt-3/2/0.32769 [VPLS/7] 1d 13:54:49
              Unusable
vt-3/2/0.32772 [VPLS/7] 1d 13:54:59
              Unusable
```

show route aspath-regex

Syntax	<code>show route aspath-regex <i>regular-expression</i></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches)	<code>show route aspath-regex <i>regular-expression</i></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the entries in the routing table that match the specified autonomous system (AS) path regular expression.
Options	<p><i>regular-expression</i>—Regular expression that matches an entire AS path.</p> <p><i>logical-system (all logical-system-name)</i>—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Additional Information	<p>You can specify a regular expression as:</p> <ul style="list-style-type: none">• An individual AS number• A period wildcard used in place of an AS number• An AS path regular expression that is enclosed in parentheses <p>You also can include the operators described in the table of AS path regular expression operators in the <i>Junos Policy Framework Configuration Guide</i>. The following list summarizes these operators:</p> <ul style="list-style-type: none">• <i>{m,n}</i>—At least <i>m</i> and at most <i>n</i> repetitions of the AS path term.• <i>{m}</i>—Exactly <i>m</i> repetitions of the AS path term.• <i>{m,}</i>—<i>m</i> or more repetitions of the AS path term.• <i>*</i>—Zero or more repetitions of an AS path term.• <i>+</i>—One or more repetitions of an AS path term.• <i>?</i>—Zero or one repetition of an AS path term.• <i>aspath_term aspath_term</i>—Match one of the two AS path terms. <p>When you specify more than one AS number or path term, or when you include an operator in the regular expression, enclose the entire regular expression in quotation marks. For example, to match any path that contains AS number 234, specify the following command:</p> <pre>show route aspath-regex ". * 234 . *"</pre>
Required Privilege Level	view

Related Documentation	<ul style="list-style-type: none"> • <i>Example: Using AS Path Regular Expressions</i>
List of Sample Output	show route aspath-regex (Matching a Specific AS Number) on page 791 show route aspath-regex (Matching Any Path with Two AS Numbers) on page 791
Output Fields	For information about output fields, see the output field table for the show route command.

Sample Output

show route aspath-regex (Matching a Specific AS Number)

```

user@host> show route aspath-regex 65477
inet.0: 46411 destinations, 46411 routes (46409 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

111.222.1.0/25      *[BGP/170] 00:08:48, localpref 100, from 111.222.2.24
                   AS Path: [65477] ({65488 65535}) IGP
                   to 111.222.18.225 via fpa0.0(111.222.18.233)
111.222.1.128/25   *[IS-IS/15] 09:15:37, metric 37, tag 1
                   to 111.222.18.225 via fpa0.0(111.222.18.233)
                   [BGP/170] 00:08:48, localpref 100, from 111.222.2.24
                   AS Path: [65477] ({65488 65535}) IGP
                   to 111.222.18.225 via fpa0.0(111.222.18.233)
...

```

show route aspath-regex (Matching Any Path with Two AS Numbers)

```

user@host> show route aspath-regex ?.* 234 3561.*?
inet.0: 46351 destinations, 46351 routes (46349 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

9.20.0.0/17        *[BGP/170] 01:35:00, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 2685 2686 Incomplete
                   to 192.156.169.1 via 192.156.169.14(so-0/0/0)
12.10.231.0/24     *[BGP/170] 01:35:00, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 5696 7369 IGP
                   to 192.156.169.1 via 192.156.169.14(so-0/0/0)
24.64.32.0/19      *[BGP/170] 01:34:59, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 6327 IGP
                   to 192.156.169.1 via 192.156.169.14(so-0/0/0)
...

```

show route best

Syntax	<code>show route best <i>destination-prefix</i></code> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	<code>show route best <i>destination-prefix</i></code> <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the route in the routing table that is the best route to the specified address or range of addresses. The best route is the longest matching route.
Options	brief detail extensive terse —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief . <i>destination-prefix</i> —Address or range of addresses. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route best on page 792 show route best detail on page 793 show route best extensive on page 794 show route best terse on page 794
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route best

```
user@host> show route best 10.255.70.103
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.255.70.103/32    *[OSPF/10] 1d 13:19:20, metric 2
                  > to 10.31.1.6 via ge-3/1/0.0
                  via so-0/3/0.0

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.255.70.103/32    *[RSVP/7] 1d 13:20:13, metric 2
                  > via so-0/3/0.0, label-switched-path green-r1-r3

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
```

```

+ = Active Route, - = Last Active, * = Both
10.0.0.0/8      *[Direct/0] 2d 01:43:34
                 > via fxp2.0
                 [Direct/0] 2d 01:43:34
                 > via fxp1.0

```

show route best detail

```

user@host> show route best 10.255.70.103 detail
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
10.255.70.103/32 (1 entry, 1 announced)
    *OSPF    Preference: 10
              Next-hop reference count: 9
              Next hop: 10.31.1.6 via ge-3/1/0.0, selected
              Next hop: via so-0/3/0.0
              State: <Active Int>
              Local AS:    69
              Age: 1d 13:20:06      Metric: 2
              Area: 0.0.0.0
              Task: OSPF
              Announcement bits (2): 0-KRT 3-Resolve tree 2
              AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
10.255.70.103/32 (1 entry, 1 announced)
    State: <FlashAll>
    *RSVP    Preference: 7
              Next-hop reference count: 5
              Next hop: via so-0/3/0.0 weight 0x1, selected
              Label-switched-path green-r1-r3
              Label operation: Push 100016
              State: <Active Int>
              Local AS:    69
              Age: 1d 13:20:59      Metric: 2
              Task: RSVP
              Announcement bits (1): 1-Resolve tree 2
              AS path: I

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
10.0.0.0/8 (2 entries, 0 announced)
    *Direct Preference: 0
              Next hop type: Interface
              Next-hop reference count: 1
              Next hop: via fxp2.0, selected
              State: <Active Int>
              Age: 2d 1:44:20
              Task: IF
              AS path: I
    Direct Preference: 0
              Next hop type: Interface
              Next-hop reference count: 1
              Next hop: via fxp1.0, selected
              State: <NotBest Int>
              Inactive reason: No difference
              Age: 2d 1:44:20
              Task: IF
              AS path: I

```

show route best extensive

The output for the **show route best extensive** command is identical to that for the **show route best detail** command. For sample output, see [show route best detail on page 793](#).

show route best terse

```
user@host> show route best 10.255.70.103 terse
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* 10.255.70.103/32  0 10      2          >10.31.1.6
                               so-0/3/0.0

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* 10.255.70.103/32  R  7      2          >so-0/3/0.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* 10.0.0.0/8        D  0          >fxp2.0
                    D  0          >fxp1.0
```

show route brief

Syntax	show route brief <destination-prefix> <logical-system (all logical-system-name)>
Syntax (EX Series Switches)	show route brief <destination-prefix>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display brief information about the active entries in the routing tables.
Options	<p>none—Display all active entries in the routing table.</p> <p>destination-prefix—(Optional) Display active entries for the specified address or range of addresses.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show route brief on page 795
Output Fields	For information about output fields, see the Output Field table of the show route command.

Sample Output

show route brief

```

user@host> show route brief
inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 1w5d 20:30:29
                   Discard
10.255.245.51/32   *[Direct/0] 2w4d 13:11:14
                   > via lo0.0
172.16.0.0/12      *[Static/5] 2w4d 13:11:14
                   > to 192.168.167.254 via fxp0.0
192.168.0.0/18     *[Static/5] 1w5d 20:30:29
                   > to 192.168.167.254 via fxp0.0
192.168.40.0/22    *[Static/5] 2w4d 13:11:14
                   > to 192.168.167.254 via fxp0.0
192.168.64.0/18    *[Static/5] 2w4d 13:11:14
                   > to 192.168.167.254 via fxp0.0
192.168.164.0/22   *[Direct/0] 2w4d 13:11:14
                   > via fxp0.0
192.168.164.51/32  *[Local/0] 2w4d 13:11:14
                   Local via fxp0.0
207.17.136.192/32 *[Static/5] 2w4d 13:11:14

```

```

> to 192.168.167.254 via fxp0.0
green.inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
100.101.0.0/16    *[Direct/0] 1w5d 20:30:28
                  > via fe-0/0/3.0
100.101.2.3/32   *[Local/0] 1w5d 20:30:28
                  Local via fe-0/0/3.0
224.0.0.5/32     *[OSPF/10] 1w5d 20:30:29, metric 1
                  MultiRecv
```

show route ccc

Syntax	show route ccc ccc <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display circuit cross-connect (CCC) entries in the Multiprotocol Link Switching (MPLS) routing table.
Options	<p>ccc—Name of an entry with a circuit cross-connect interface.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>show connections</i>
List of Sample Output	show route ccc extensive on page 797
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route ccc extensive

```

user@host> show route ccc fe-0/1/0.600 extensive
mpls.0: 19 destinations, 19 routes (19 active, 0 holddown, 0 hidden)
fe-0/1/2.600 (1 entry, 1 announced)
TSI:
KRT in-kernel fe-0/1/2.600.0      /16 -> {0.0.0.0}
  *CCC      Preference: 7
            Next-hop reference count: 2
            Next hop: via so-0/0/3.0 weight 0x1, selected
            Label operation: Push 101424
            State: <Active Int>
            Local AS: 100
            Age: 28:13   Metric: 3
            Task: MPLS
            Announcement bits (1): 0-KRT
            AS path: I

```

show route community

Syntax	<code>show route community <i>as-number:community-value</i></code> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	<code>show route community <i>as-number:community-value</i></code> <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the route entries in each routing table that are members of a Border Gateway Protocol (BGP) community.
Options	<p><i>as-number:community-value</i>—One or more community identifiers. <i>as-number</i> is the AS number, and <i>community-value</i> is the community identifier. When you specify more than one community identifier, enclose the identifiers in double quotation marks. Community identifiers can include wildcards.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Additional Information	Specifying the community option displays all routes matching the community found within the routing table. The community option does not limit the output to only the routes being advertised to the neighbor after any egress routing policy.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• show route detail on page 807
List of Sample Output	show route community on page 798
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route community

```
user@host> show route community 234:80
inet.0: 46511 destinations, 46511 routes (46509 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

4.0.0.0/8          *[BGP/170] 03:33:07, localpref 100, from 131.103.20.49
                   AS Path: {666} 234 2548 1 IGP
                   to 192.156.169.1 via 192.156.169.14(so-0/0/0)
6.0.0.0/8          *[BGP/170] 03:33:07, localpref 100, from 131.103.20.49
```



```
9.2.0.0/16      AS Path: {666} 234 2548 568 721 Incomplete
                  to 192.156.169.1 via 192.156.169.14(so-0/0/0)
                  *[BGP/170] 03:33:06, localpref 100, from 131.103.20.49
                  AS Path: {666} 234 2548 1673 1675 1747 IGP
                  to 192.156.169.1 via 192.156.169.14(so-0/0/0)
```

show route community-name

Syntax	<code>show route community-name <i>community-name</i></code> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	<code>show route community-name <i>community-name</i></code> <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the route entries in each routing table that are members of a Border Gateway Protocol (BGP) community, specified by a community name.
Options	<i>community-name</i> —Name of the community. brief detail extensive terse —(Optional) Display the specified level of output. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route community-name on page 800
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route community-name

```
user@host> show route community-name red-com
inet.0: 17 destinations, 17 routes (16 active, 0 holddown, 1 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

instance1.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.245.212/32  *[BGP/170] 00:04:40, localpref 100, from 10.255.245.204
                  AS path: 300 I
                  > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
20.20.20.20/32    *[BGP/170] 00:04:40, localpref 100, from 10.255.245.204
                  AS path: I
                  > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
100.1.4.0/24     *[BGP/170] 00:04:40, localpref 100, from 10.255.245.204
                  AS path: I
                  > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
```

```
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.245.204:10:10.255.245.212/32
    *[BGP/170] 00:06:40, localpref 100, from 10.255.245.204
        AS path: 300 I
        > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
10.255.245.204:10:20.20.20.20/32
    *[BGP/170] 00:36:02, localpref 100, from 10.255.245.204
        AS path: I
        > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
10.255.245.204:10:100.1.4.0/24
    *[BGP/170] 00:36:02, localpref 100, from 10.255.245.204
        AS path: I
        > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

instance1.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route damping

Syntax	show route damping (decayed history suppressed) <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>	
Syntax (EX Series Switch and QFX Series)	show route damping (decayed history suppressed) <brief detail extensive terse>	
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.3 for the QFX Series.	
Description	Display the BGP routes for which updates might have been reduced because of route flap damping.	
Options	<p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>decayed—Display route damping entries that might no longer be valid, but are not suppressed.</p> <p>history—Display entries that have already been withdrawn, but have been logged.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>suppressed—Display entries that have been suppressed and are no longer being installed into the forwarding table or exported by routing protocols.</p>	
Required Privilege Level	view	
Related Documentation	<ul style="list-style-type: none"> • clear bgp damping on page 308 • show policy damping on page 347 	
List of Sample Output	show route damping decayed detail on page 805 show route damping history on page 806 show route damping history detail on page 806	
Output Fields	Table 125 on page 802 lists the output fields for the show route damping command. Output fields are listed in the approximate order in which they appear.	

Table 125: show route damping Output Fields

Field Name	Field Description	Level of Output
<i>routing-table-name</i>	Name of the routing table—for example, <i>inet.0</i> .	All levels
<i>destinations</i>	Number of destinations for which there are routes in the routing table.	All levels

Table 125: show route damping Output Fields (*continued*)

Field Name	Field Description	Level of Output
<i>number routes</i>	Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> • active • holddown (routes that are in a pending state before being declared inactive) • hidden (the routes are not used because of a routing policy) 	All levels
<i>destination-prefix (entry, announced)</i>	Destination prefix. The entry value is the number of routes for this destination, and the announced value is the number of routes being announced for this destination.	detail extensive
<i>[protocol, preference]</i>	Protocol from which the route was learned and the preference value for the route. <ul style="list-style-type: none"> • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • -—A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route. <p>In every routing metric except for the BGP LocalPref attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the LocalPref value in the Preference2 field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the LocalPref value for Route 2 is 155, the Preference2 value is -156. Route 2 is preferred because it has a higher LocalPref value and a lower Preference2 value.</p>	All levels
Next-hop reference count	Number of references made to the next hop.	detail extensive
Source	IP address of the route source.	detail extensive
Next hop	Network layer address of the directly reachable neighboring system.	detail extensive
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word Selected .	detail extensive
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.	detail extensive
Indirect next hop	Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.	detail extensive
State	Flags for this route. For a description of possible values for this field, see the output field table for the show route detail command.	detail extensive
Local AS	AS number of the local routing device.	detail extensive
Peer AS	AS number of the peer routing device.	detail extensive

Table 125: show route damping Output Fields (*continued*)

Field Name	Field Description	Level of Output
Age	How long the route has been known.	detail extensive
Metric	Metric for the route.	detail extensive
Task	Name of the protocol that has added the route.	detail extensive
Announcement bits	List of protocols that announce this route. n-Resolve inet indicates that the route is used for route resolution for next hops found in the routing table. n is an index used by Juniper Networks customer support only.	detail extensive
AS path	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • ?—Incomplete; typically, the AS path was aggregated. <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> • []—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device or if AS path prepending is configured. • { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • ()—Parentheses enclose a confederation. • ([])—Parentheses and brackets enclose a confederation set. <p>NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p>	All levels
to	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.	brief none
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word Selected .	brief none
Communities	Community path attribute for the route. See the output field table for the show route detail command.	detail extensive
Localpref	Local preference value included in the route.	All levels
Router ID	BGP router ID as advertised by the neighbor in the open message.	detail extensive
Merit (last update/now)	Last updated and current figure-of-merit value.	detail extensive

Table 125: show route damping Output Fields (*continued*)

Field Name	Field Description	Level of Output
damping-parameters	Name that identifies the damping parameters used, which is defined in the damping statement at the [edit policy-options] hierarchy level.	detail extensive
Last update	Time of most recent change in path attributes.	detail extensive
First update	Time of first change in path attributes, which started the route damping process.	detail extensive
Flaps	Number of times the route has gone up or down or its path attributes have changed.	detail extensive
Suppressed	(suppressed keyword only) This route is currently suppressed. A suppressed route does not appear in the forwarding table and routing protocols do not export it.	All levels
Reusable in	(suppressed keyword only) Time when a suppressed route will again be available.	All levels
Preference will be	(suppressed keyword only) Preference value that will be applied to the route when it is again active.	All levels

Sample Output

show route damping decayed detail

```

user@host> show route damping decayed detail
inet.0: 173319 destinations, 1533668 routes (172625 active, 4 holddown, 108083
hidden)
10.0.111.0/24 (7 entries, 1 announced)
  *BGP      Preference: 170/-101
            Next-hop reference count: 151973
            Source: 172.23.2.129
            Next hop: via so-1/2/0.0
            Next hop: via so-5/1/0.0, selected
            Next hop: via so-6/0/0.0
            Protocol next hop: 172.23.2.129
            Indirect next hop: 89a1a00 264185
            State: <Active Ext>
            Local AS: 65000 Peer AS: 65490
            Age: 3:28      Metric2: 0
            Task: BGP_65490.172.23.2.129+179
            Announcement bits (6): 0-KRT 1-RT 4-KRT 5-BGP.0.0.0.0+179

6-Resolve tree 2 7-Resolve tree 3
AS path: 65490 65520 65525 65525 65525 65525 I ()
Communities: 65501:390 65501:2000 65501:3000 65504:701
Localpref: 100
Router ID: 172.23.2.129
Merit (last update/now): 1934/1790
damping-parameters: damping-high
Last update:          00:03:28 First update:          00:06:40
Flaps: 2

```

show route damping history

```
user@host> show route damping history
inet.0: 173320 destinations, 1533529 routes (172624 active, 6 holddown, 108122
hidden)
+ = Active Route, - = Last Active, * = Both

10.108.0.0/15      [BGP ] 2d 22:47:58, localpref 100
                  AS path: 65220 65501 65502 I
                  > to 192.168.60.85 via so-3/1/0.0
```

show route damping history detail

```
user@host> show route damping history detail
inet.0: 173319 destinations, 1533435 routes (172627 active, 2 holddown, 108105
hidden)
10.108.0.0/15 (3 entries, 1 announced)
    BGP                /-101
        Next-hop reference count: 69058
        Source: 192.168.60.85
        Next hop: 192.168.60.85 via so-3/1/0.0, selected
        State: <Hidden Ext>
        Inactive reason: Unusable path
        Local AS: 65000 Peer AS: 65220
        Age: 2d 22:48:10
        Task: BGP_65220.192.168.60.85+179
        AS path: 65220 65501 65502 I ()
        Communities: 65501:390 65501:2000 65501:3000 65504:3561
        Localpref: 100
        Router ID: 192.168.80.25
        Merit (last update/now): 1000/932
        damping-parameters: set-normal
        Last update:          00:01:05 First update:          00:01:05
        Flaps: 1
```


show route detail

Syntax	show route detail <destination-prefix> <logical-system (all logical-system-name)>
Syntax (EX Series Switches)	show route detail <destination-prefix>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Description	Display detailed information about the active entries in the routing tables.
Options	<p>none—Display all active entries in the routing table on all systems.</p> <p>destination-prefix—(Optional) Display active entries for the specified address or range of addresses.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show route detail on page 816 show route detail (with BGP Multipath) on page 822 show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 822 show route label detail (Multipoint LDP with Multicast-Only Fast Reroute) on page 823
Output Fields	<p>Table 126 on page 807 describes the output fields for the show route detail command. Output fields are listed in the approximate order in which they appear.</p>

Table 126: show route detail Output Fields

Field Name	Field Description
<i>routing-table-name</i>	Name of the routing table (for example, inet.0).
<i>number destinations</i>	Number of destinations for which there are routes in the routing table.
<i>number routes</i>	Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> active (routes that are active) holddown (routes that are in the pending state before being declared inactive) hidden (routes that are not used because of a routing policy)

Table 126: show route detail Output Fields (*continued*)

Field Name	Field Description
<i>route-destination</i> (entry, announced)	<p>Route destination (for example:10.0.0.1/24). The entry value is the number of routes for this destination, and the announced value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:</p> <ul style="list-style-type: none"> • MPLS-label (for example, 80001). • interface-name (for example, ge-1/0/2). • neighbor-address:control-word-status:encapsulation type:vc-id:source (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96). <ul style="list-style-type: none"> • neighbor-address—Address of the neighbor. • control-word-status—Whether the use of the control word has been negotiated for this virtual circuit: NoCtrlWord or CtrlWord. • encapsulation type—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport. • vc-id—Virtual circuit identifier. • source—Source of the advertisement: Local or Remote.
label stacking	<p>(Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels.</p> <ul style="list-style-type: none"> • S=0 route indicates that a packet with an incoming label stack depth of 2 or more exits this routing device with one fewer label (the label-popping operation is performed). • If there is no S= information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).
[<i>protocol, preference</i>]	<p>Protocol from which the route was learned and the preference value for the route.</p> <ul style="list-style-type: none"> • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • - —A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route. <p>In every routing metric except for the BGP LocalPref attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the LocalPref value in the Preference2 field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the LocalPref value for Route 2 is 155, the Preference2 value is -156. Route 2 is preferred because it has a higher LocalPref value and a lower Preference2 value.</p>
Level	<p>(IS-IS only). In IS-IS, a single AS can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.</p>
Route Distinguisher	IP subnet augmented with a 64-bit prefix.
PMSI	Provider multicast service interface (MVPN routing table).
Next-hop type	Type of next hop. For a description of possible values for this field, see Table 127 on page 812 .

Table 126: show route detail Output Fields (*continued*)

Field Name	Field Description
Next-hop reference count	Number of references made to the next hop.
Flood nexthop branches exceed maximum message	Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.
Source	IP address of the route source.
Next hop	Network layer address of the directly reachable neighboring system.
via	<p>Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word Selected. This field can also contain the following information:</p> <ul style="list-style-type: none"> • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.
Label-switched-path lsp-path-name	Name of the LSP used to reach the next hop.
Label operation	MPLS label and operation occurring at this routing device. The operation can be pop (where a label is removed from the top of the stack), push (where another label is added to the label stack), or swap (where a label is replaced by another label).
Interface	(Local only) Local interface name.
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.
Indirect next hop	Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.
State	State of the route (a route can be in more than one state). See Table 128 on page 813 .
Local AS	AS number of the local routing device.
Age	How long the route has been known.
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.
Metricn	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.

Table 126: show route detail Output Fields (*continued*)

Field Name	Field Description
MED-plus-IGP	Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.
TTL-Action	<p>For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances.</p> <p>For sample output, see show route table.</p>
Task	Name of the protocol that has added the route.
Announcement bits	List of protocols that announce this route. n-Resolve inet indicates that the route is used for route resolution for next hops found in the routing table. n is an index used by Juniper Networks customer support only.
AS path	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • Recorded—The AS path is recorded by the sample process (sampled). • ?—Incomplete; typically, the AS path was aggregated. <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> • []—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used in the AS-path merge process, as defined in RFC 4893. • []—If more than one AS number is configured on the routing device, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path. • { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • ()—Parentheses enclose a confederation. • ([])—Parentheses and brackets enclose a confederation set. <p>NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p>
validation-state	<p>(BGP-learned routes) Validation status of the route:</p> <ul style="list-style-type: none"> • Invalid—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • Unknown—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • Unverified—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers. • Valid—Indicates that the prefix and autonomous system pair are found in the database.
FECs bound to route	Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.

Table 126: show route detail Output Fields (*continued*)

Field Name	Field Description
Primary Upstream	When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path.
RPF Nexthops	When multipoint LDP with MoFRR is configured, the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.
Label	Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.
weight	Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.
VC Label	MPLS label assigned to the Layer 2 circuit virtual connection.
MTU	Maximum transmission unit (MTU) of the Layer 2 circuit.
VLAN ID	VLAN identifier of the Layer 2 circuit.
Prefixes bound to route	Forwarding equivalent class (FEC) bound to this route. Applicable only to routes installed by LDP.
Communities	Community path attribute for the route. See Table 129 on page 815 for all possible values for this field.
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).
control flags	Control flags: none or Site Down .
mtu	Maximum transmission unit (MTU) information.
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.
status vector	Layer 2 VPN and VPLS network layer reachability information (NLRI).
Accepted Multipath	Current active path when BGP multipath is configured.
Accepted MultipathContrib	Path currently contributing to BGP multipath.
Localpref	Local preference value included in the route.
Router ID	BGP router ID as advertised by the neighbor in the open message.
Primary Routing Table	In a routing table group, the name of the primary routing table in which the route resides.
Secondary Tables	In a routing table group, the name of one or more secondary tables in which the route resides.

Table 127 on page 812 describes all possible values for the Next-hop Types output field.

Table 127: Next-hop Types Output Field Values

Next-Hop Type	Description
Broadcast (bcast)	Broadcast next hop.
Deny	Deny next hop.
Discard	Discard next hop.
Flood	Flood next hop. Consists of components called branches, up to a maximum of 32 branches. Each flood next-hop branch sends a copy of the traffic to the forwarding interface. Used by point-to-multipoint RSVP, point-to-multipoint LDP, point-to-multipoint CCC, and multicast.
Hold	Next hop is waiting to be resolved into a unicast or multicast type.
Indexed (idxd)	Indexed next hop.
Indirect (indr)	Used with applications that have a protocol next hop address that is remote. You are likely to see this next-hop type for internal BGP (IBGP) routes when the BGP next hop is a BGP neighbor that is not directly connected.
Interface	Used for a network address assigned to an interface. Unlike the router next hop, the interface next hop does not reference any specific node on the network.
Local (locl)	Local address on an interface. This next-hop type causes packets with this destination address to be received locally.
Multicast (mcst)	Wire multicast next hop (limited to the LAN).
Multicast discard (mdsc)	Multicast discard.
Multicast group (mgrp)	Multicast group member.
Receive (recv)	Receive.
Reject (rjct)	Discard. An ICMP unreachable message was sent.
Resolve (rslv)	Resolving next hop.
Routed multicast (mcrt)	Regular multicast next hop.

Table 127: Next-hop Types Output Field Values (*continued*)

Next-Hop Type	Description
Router	<p>A specific node or set of nodes to which the routing device forwards packets that match the route prefix.</p> <p>To qualify as next-hop type router, the route must meet the following criteria:</p> <ul style="list-style-type: none"> • Must not be a direct or local subnet for the routing device. • Must have a next hop that is directly connected to the routing device.
Table	Routing table next hop.
Unicast (ucst)	Unicast.
Unilist (ulst)	List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.

Table 128 on page 813 describes all possible values for the State output field. A route can be in more than one state (for example, <Active NoReadvrt Int Ext>).

Table 128: State Output Field Values

Value	Description
Accounting	Route needs accounting.
Active	Route is active.
Always Compare MED	Path with a lower multiple exit discriminator (MED) is available.
AS path	Shorter AS path is available.
Cisco Non-deterministic MED selection	Cisco nondeterministic MED is enabled, and a path with a lower MED is available.
Clone	Route is a clone.
Cluster list length	Length of cluster list sent by the route reflector.
Delete	Route has been deleted.
Ex	Exterior route.
Ext	BGP route received from an external BGP neighbor.

Table 128: State Output Field Values (*continued*)

Value	Description
FlashAll	Forces all protocols to be notified of a change to any route, active or inactive, for a prefix. When not set, protocols are informed of a prefix only when the active route changes.
Hidden	Route not used because of routing policy.
IfCheck	Route needs forwarding RPF check.
IGP metric	Path through next hop with lower IGP metric is available.
Inactive reason	Flags for this route, which was not selected as best for a particular destination.
Initial	Route being added.
Int	Interior route.
Int Ext	BGP route received from an internal BGP peer or a BGP confederation peer.
Interior > Exterior > Exterior via Interior	Direct, static, IGP, or EBGP path is available.
Local Preference	Path with a higher local preference value is available.
Martian	Route is a martian (ignored because it is obviously invalid).
MartianOK	Route exempt from martian filtering.
Next hop address	Path with lower metric next hop is available.
No difference	Path from neighbor with lower IP address is available.
NoReadvrt	Route not to be advertised.
NotBest	Route not chosen because it does not have the lowest MED.
Not Best in its group	Incoming BGP AS is not the best of a group (only one AS can be the best).
NotInstall	Route not to be installed in the forwarding table.
Number of gateways	Path with a greater number of next hops is available.
Origin	Path with a lower origin code is available.
Pending	Route pending because of a hold-down configured on another route.

Table 128: State Output Field Values (*continued*)

Value	Description
Release	Route scheduled for release.
RIB preference	Route from a higher-numbered routing table is available.
Route Distinguisher	64-bit prefix added to IP subnets to make them unique.
Route Metric or MED comparison	Route with a lower metric or MED is available.
Route Preference	Route with lower preference value is available
Router ID	Path through a neighbor with lower ID is available.
Secondary	Route not a primary route.
Unusable path	Path is not usable because of one of the following conditions: <ul style="list-style-type: none"> • The route is damped. • The route is rejected by an import policy. • The route is unresolved.
Update source	Last tiebreaker is the lowest IP address value.

Table 129 on page 815 describes the possible values for the Communities output field.

Table 129: Communities Output Field Values

Value	Description
<i>area-number</i>	4 bytes, encoding a 32-bit area number. For AS-external routes, the value is 0. A nonzero value identifies the route as internal to the OSPF domain, and as within the identified area. Area numbers are relative to a particular OSPF domain.
bandwidth: local AS number:link-bandwidth-number	Link-bandwidth community value used for unequal-cost load balancing. When BGP has several candidate paths available for multipath purposes, it does not perform unequal-cost load balancing according to the link-bandwidth community unless all candidate paths have this attribute.
domain-id	Unique configurable number that identifies the OSPF domain.
domain-id-vendor	Unique configurable number that further identifies the OSPF domain.
<i>link-bandwidth-number</i>	Link-bandwidth number: from 0 through 4,294,967,295 (bytes per second).
<i>local AS number</i>	Local AS number: from 1 through 65,535.
<i>options</i>	1 byte. Currently this is only used if the route type is 5 or 7. Setting the least significant bit in the field indicates that the route carries a type 2 metric.

Table 129: Communities Output Field Values (*continued*)

Value	Description
origin	(Used with VPNs) Identifies where the route came from.
ospf-route-type	1 byte, encoded as 1 or 2 for intra-area routes (depending on whether the route came from a type 1 or a type 2 LSA); 3 for summary routes; 5 for external routes (area number must be 0); 7 for NSSA routes; or 129 for sham link endpoint addresses.
route-type-vendor	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute 0x8000 . The format is area-number:ospf-route-type:options .
rte-type	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute 0x0306 . The format is area-number:ospf-route-type:options .
target	Defines which VPN the route participates in; target has the format 32-bit IP address:16-bit number . For example, 10.19.0.0:100.
unknown IANA	Incoming IANA codes with a value between 0x1 and 0x7fff . This code of the BGP extended community attribute is accepted, but it is not recognized.
unknown OSPF vendor community	Incoming IANA codes with a value above 0x8000 . This code of the BGP extended community attribute is accepted, but it is not recognized.

Sample Output

show route detail

```

user@host> show route detail

inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 29
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 1:31:43
    Task: RT
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I

10.31.1.0/30 (2 entries, 1 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 2
    Next hop: via so-0/3/0.0, selected
    State: <Active Int>
    Local AS: 69
    Age: 1:30:17
    Task: IF
    Announcement bits (1): 3-Resolve tree 2
    AS path: I
  OSPF Preference: 10

```

```

Next-hop reference count: 1
Next hop: via so-0/3/0.0, selected
State: <Int>
Inactive reason: Route Preference
Local AS: 69
Age: 1:30:17 Metric: 1
Area: 0.0.0.0
Task: OSPF
AS path: I

10.31.1.1/32 (1 entry, 1 announced)
  *Local Preference: 0
    Next hop type: Local
    Next-hop reference count: 7
    Interface: so-0/3/0.0
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:30:20
    Task: IF
    Announcement bits (1): 3-Resolve tree 2
    AS path: I

...

10.31.2.0/30 (1 entry, 1 announced)
  *OSPF Preference: 10
    Next-hop reference count: 9
    Next hop: via so-0/3/0.0
    Next hop: 10.31.1.6 via ge-3/1/0.0, selected
    State: <Active Int>
    Local AS: 69
    Age: 1:29:56 Metric: 2
    Area: 0.0.0.0
    Task: OSPF
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I

...

224.0.0.2/32 (1 entry, 1 announced)
  *PIM Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:31:45
    Task: PIM Recv
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I

...

224.0.0.22/32 (1 entry, 1 announced)
  *IGMP Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:31:43
    Task: IGMP
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I

```

```
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.255.70.103/32 (1 entry, 1 announced)
  State: <FlashAll>
  *RSVP Preference: 7
    Next-hop reference count: 6
    Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
    Label-switched-path green-r1-r3
    Label operation: Push 100096
    State: <Active Int>
    Local AS: 69
    Age: 1:25:49 Metric: 2
    Task: RSVP
    Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
    AS path: I

10.255.71.238/32 (1 entry, 1 announced)
  State: <FlashAll>
  *RSVP Preference: 7
    Next-hop reference count: 6
    Next hop: via so-0/3/0.0 weight 0x1, selected
    Label-switched-path green-r1-r2
    State: <Active Int>
    Local AS: 69
    Age: 1:25:49 Metric: 1
    Task: RSVP
    Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
    AS path: I

private__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Active Int>
    Local AS: 69
    Age: 1:31:44
    Task: IF
    AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
0 (1 entry, 1 announced)
  *MPLS Preference: 0
    Next hop type: Receive
    Next-hop reference count: 6
    State: <Active Int>
    Local AS: 69
    Age: 1:31:45 Metric: 1
    Task: MPLS
    Announcement bits (1): 0-KRT
    AS path: I

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

299840 (1 entry, 1 announced)
```

```

TSI:
KRT in-kerne 299840 /52 -> {indirect(1048575)}
    *RSVP Preference: 7/2
        Next hop type: Flood
        Address: 0x9174a30
        Next-hop reference count: 4
        Next hop type: Router, Next hop index: 798
        Address: 0x9174c28
        Next-hop reference count: 2
        Next hop: 8.0.0.2 via lt-1/2/0.9 weight 0x1
        Label-switched-path R2-to-R4-2p2mp
        Label operation: Pop
        Next hop type: Router, Next hop index: 1048574
        Address: 0x92544f0
        Next-hop reference count: 2
        Next hop: 7.0.0.2 via lt-1/2/0.7 weight 0x1
        Label-switched-path R2-to-R200-p2mp
        Label operation: Pop
        Next hop: 6.0.0.2 via lt-1/2/0.5 weight 0x8001
        Label operation: Pop
        State: <Active Int>
        Age: 1:29 Metric: 1
        Task: RSVP
        Announcement bits (1): 0-KRT
        AS path: I...

800010 (1 entry, 1 announced)
    *VPLS Preference: 7
        Next-hop reference count: 2
        Next hop: via vt-3/2/0.32769, selected
        Label operation: Pop
        State: <Active Int>
        Age: 1:29:30
        Task: Common L2 VC
        Announcement bits (1): 0-KRT
        AS path: I

vt-3/2/0.32769 (1 entry, 1 announced)
    *VPLS Preference: 7
        Next-hop reference count: 2
        Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
        Label-switched-path green-r1-r3
        Label operation: Push 800012, Push 100096(top)
        Protocol next hop: 10.255.70.103
        Push 800012
        Indirect next hop: 87272e4 1048574
        State: <Active Int>
        Age: 1:29:30 Metric2: 2
        Task: Common L2 VC
        Announcement bits (2): 0-KRT 1-Common L2 VC
        AS path: I
        Communities: target:11111:1 Layer2-info: encaps:VPLS,
        control flags:, mtu: 0

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

abcd::10:255:71:52/128 (1 entry, 0 announced)
    *Direct Preference: 0
        Next hop type: Interface
        Next-hop reference count: 1
        Next hop: via lo0.0, selected

```

```
State: <Active Int>
Local AS: 69
Age: 1:31:44
Task: IF
AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 1
Next hop: via lo0.0, selected
State: <Active NoReadvrt Int>
Local AS: 69
Age: 1:31:44
Task: IF
AS path: I

ff02::2/128 (1 entry, 1 announced)
*PIM Preference: 0
Next-hop reference count: 18
State: <Active NoReadvrt Int>
Local AS: 69
Age: 1:31:45
Task: PIM Recv6
Announcement bits (1): 0-KRT
AS path: I

ff02::d/128 (1 entry, 1 announced)
*PIM Preference: 0
Next-hop reference count: 18
State: <Active NoReadvrt Int>
Local AS: 69
Age: 1:31:45
Task: PIM Recv6
Announcement bits (1): 0-KRT
AS path: I

ff02::16/128 (1 entry, 1 announced)
*MLD Preference: 0
Next-hop reference count: 18
State: <Active NoReadvrt Int>
Local AS: 69
Age: 1:31:43
Task: MLD
Announcement bits (1): 0-KRT
AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
*Direct Preference: 0
Next hop type: Interface
Next-hop reference count: 1
Next hop: via lo0.16385, selected
State: <Active NoReadvrt Int>
Age: 1:31:44
Task: IF
AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
```

```

10.255.70.103:1:3:1/96 (1 entry, 1 announced)
  *BGP   Preference: 170/-101
        Route Distinguisher: 10.255.70.103:1
        Next-hop reference count: 7
        Source: 10.255.70.103
        Protocol next hop: 10.255.70.103
        Indirect next hop: 2 no-forward
        State: <Secondary Active Int Ext>
        Local AS: 69 Peer AS: 69
        Age: 1:25:49 Metric2: 1
        AIGP 210
        Task: BGP_69.10.255.70.103+179
        Announcement bits (1): 0-green-l2vpn
        AS path: I
        Communities: target:11111:1 Layer2-info: encaps:VPLS,
        control flags:, mtu: 0
        Label-base: 800008, range: 8
        Localpref: 100
        Router ID: 10.255.70.103
        Primary Routing Table bgp.l2vpn.0

10.255.71.52:1:1:1/96 (1 entry, 1 announced)
  *L2VPN Preference: 170/-1
        Next-hop reference count: 5
        Protocol next hop: 10.255.71.52
        Indirect next hop: 0 -
        State: <Active Int Ext>
        Age: 1:31:40 Metric2: 1
        Task: green-l2vpn
        Announcement bits (1): 1-BGP.0.0.0.0+179
        AS path: I
        Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
        mtu: 0
        Label-base: 800016, range: 8, status-vector: 0x9F

10.255.71.52:1:5:1/96 (1 entry, 1 announced)
  *L2VPN Preference: 170/-101
        Next-hop reference count: 5
        Protocol next hop: 10.255.71.52
        Indirect next hop: 0 -
        State: <Active Int Ext>
        Age: 1:31:40 Metric2: 1
        Task: green-l2vpn
        Announcement bits (1): 1-BGP.0.0.0.0+179
        AS path: I
        Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
        Label-base: 800008, range: 8, status-vector: 0x9F

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
  *L2CKT Preference: 7
        Next hop: via so-1/1/2.0 weight 1, selected
        Label-switched-path my-lsp
        Label operation: Push 100000[0]
        Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
        State: <Active Int>
        Local AS: 99
        Age: 10:21
        Task: l2 circuit

```

```
Announcement bits (1): 0-LDP
AS path: I
VC Label 100000, MTU 1500, VLAN ID 512
```

show route detail (with BGP Multipath)

```
user@host> show route detail

10.1.1.8/30 (2 entries, 1 announced)
  *BGP   Preference: 170/-101
        Next hop type: Router, Next hop index: 262142
        Address: 0x901a010
        Next-hop reference count: 2
        Source: 10.1.1.2
        Next hop: 10.1.1.2 via ge-0/3/0.1, selected
        Next hop: 10.1.1.6 via ge-0/3/0.5
        State: <Active Ext>
        Local AS:      1 Peer AS:      2
        Age: 5:04:43
        Validation State: unverified
        Task: BGP_2.10.1.1.2+59955
        Announcement bits (1): 0-KRT
        AS path: 2 I
        Accepted Multipath
        Localpref: 100
        Router ID: 1.1.1.2
  BGP   Preference: 170/-101
        Next hop type: Router, Next hop index: 678
        Address: 0x8f97520
        Next-hop reference count: 9
        Source: 10.1.1.6
        Next hop: 10.1.1.6 via ge-0/3/0.5, selected
        State: <NotBest Ext>
        Inactive reason: Not Best in its group - Active preferred
        Local AS:      1 Peer AS:      2
        Age: 5:04:43
        Validation State: unverified
        Task: BGP_2.10.1.1.6+58198
        AS path: 2 I
        Accepted MultipathContrib
        Localpref: 100
        Router ID: 1.1.1.3
```

show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show route label 299872 detail
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
  *LDP   Preference: 9
        Next hop type: Flood
        Next-hop reference count: 3
        Address: 0x9097d90
        Next hop: via vt-0/1/0.1
        Next-hop index: 661
        Label operation: Pop
        Address: 0x9172130
        Next hop: via so-0/0/3.0
        Next-hop index: 654
        Label operation: Swap 299872
        State: **Active Int>
        Local AS: 1001
```



```

Age: 8:20      Metric: 1
Task: LDP
Announcement bits (1): 0-KRT
AS path: I
FECs bound to route: P2MP root-addr 10.255.72.166, grp 232.1.1.1,
src 192.168.142.2

```

show route label detail (Multipoint LDP with Multicast-Only Fast Reroute)

```
user@host> show route label 301568 detail
```

```

mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
  *LDP    Preference: 9
    Next hop type: Flood
    Address: 0x2735208
    Next-hop reference count: 3
    Next hop type: Router, Next hop index: 1397
    Address: 0x2735d2c
    Next-hop reference count: 3
    Next hop: 1.3.8.2 via ge-1/2/22.0
    Label operation: Pop
    Load balance label: None;
    Next hop type: Router, Next hop index: 1395
    Address: 0x2736290
    Next-hop reference count: 3
    Next hop: 1.3.4.2 via ge-1/2/18.0
    Label operation: Pop
    Load balance label: None;
    State: <Active Int AckRequest MulticastRPF>
    Local AS: 10
    Age: 54:05      Metric: 1
    Validation State: unverified
    Task: LDP
    Announcement bits (1): 0-KRT
    AS path: I
    FECs bound to route: P2MP root-addr 1.1.1.1, grp: 232.1.1.1, src:
192.168.219.11
      Primary Upstream : 1.1.1.3:0--1.1.1.2:0
        RPF Nexthops :
          ge-1/2/15.0, 1.2.94.1, Label: 301568, weight: 0x1
          ge-1/2/14.0, 1.2.3.1, Label: 301568, weight: 0x1
      Backup Upstream : 1.1.1.3:0--1.1.1.6:0
        RPF Nexthops :
          ge-1/2/20.0, 1.2.96.1, Label: 301584, weight: 0xffffe
          ge-1/2/19.0, 1.3.6.1, Label: 301584, weight: 0xffffe

```

show route exact

Syntax	<code>show route exact <i>destination-prefix</i></code> <code><brief detail extensive terse></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches)	<code>show route exact <i>destination-prefix</i></code> <code><brief detail extensive terse></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display only the routes that exactly match the specified address or range of addresses.
Options	brief detail extensive terse —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief . <i>destination-prefix</i> —Address or range of addresses. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route exact on page 824 show route exact detail on page 824 show route exact extensive on page 825 show route exact terse on page 825
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route exact

```
user@host> show route exact 207.17.136.0/24

inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
207.17.136.0/24    *[Static/5] 2d 03:30:22
                  > to 192.168.71.254 via fxp0.0
```

show route exact detail

```
user@host> show route exact 207.17.136.0/24 detail

inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
207.17.136.0/24 (1 entry, 1 announced)
    *Static Preference: 5
```

```

Next-hop reference count: 29
Next hop: 192.168.71.254 via fxp0.0, selected
State: <Active NoReadvrt Int Ext>
Local AS: 69
Age: 2d 3:30:26
Task: RT
Announcement bits (2): 0-KRT 3-Resolve tree 2
AS path: I

```

show route exact extensive

```

user@host> show route exact 207.17.136.0/24 extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
207.17.136.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.0/24 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 29
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 1:25:18
    Task: RT
    Announcement bits (2): 0-KRT 3-Resolve tree 2
    AS path: I

```

show route exact terse

```

user@host> show route exact 207.17.136.0/24 terse

inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
A Destination      P Prf  Metric 1   Metric 2   Next hop      AS path
* 207.17.136.0/24  S   5                      >192.168.71.254

```

show route export

Syntax	show route export <brief detail> <instance <instance-name> routing-table-name> <logical-system (all logical-system-name)>
Syntax (EX Series Switches)	show route export <brief detail> <instance <instance-name> routing-table-name>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display policy-based route export information. Policy-based export simplifies the process of exchanging route information between routing instances.
Options	<p>none—(Same as brief.) Display standard information about policy-based export for all instances and routing tables on all systems.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>instance <instance-name>—(Optional) Display a particular routing instance for which policy-based export is currently enabled.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>routing-table-name—(Optional) Display information about policy-based export for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the show route export inet command).</p>
Required Privilege Level	view
List of Sample Output	show route export on page 827 show route export detail on page 827 show route export instance detail on page 827
Output Fields	Table 130 on page 826 lists the output fields for the show route export command. Output fields are listed in the approximate order in which they appear.

Table 130: show route export Output Fields

Field Name	Field Description	Level of Output
Table or <i>table-name</i>	Name of the routing tables that either import or export routes.	All levels
Routes	Number of routes exported from this table into other tables. If a particular route is exported to different tables, the counter will only increment by one.	brief none
Export	Whether the table is currently exporting routes to other tables: Y or N (Yes or No).	brief none

Table 130: show route export Output Fields (*continued*)

Field Name	Field Description	Level of Output
Import	Tables currently importing routes from the originator table. (Not displayed for tables that are not exporting any routes.)	detail
Flags	(instance keyword only) Flags for this feature on this instance: <ul style="list-style-type: none"> config auto-policy—The policy was deduced from the configured IGP export policies. cleanup—Configuration information for this instance is no longer valid. config—The instance was explicitly configured. 	detail
Options	(instance keyword only) Configured option displays the type of routing tables the feature handles: <ul style="list-style-type: none"> unicast—Indicates <i>instance.inet.0</i>. multicast—Indicates <i>instance.inet.2</i>. unicast multicast—Indicates <i>instance.inet.0</i> and <i>instance.inet.2</i>. 	detail
Import policy	(instance keyword only) Policy that route export uses to construct the import-export matrix. Not displayed if the instance type is vrf .	detail
Instance	(instance keyword only) Name of the routing instance.	detail
Type	(instance keyword only) Type of routing instance: forwarding , non-forwarding , or vrf .	detail

Sample Output

show route export

```

user@host> show route export
Table      Export      Routes
inet.0     N            0
black.inet.0 Y           3
red.inet.0 Y            4

```

show route export detail

```

user@host> show route export detail
inet.0                                Routes:    0
black.inet.0                          Routes:    3
  Import: [ inet.0 ]
red.inet.0                            Routes:    4
  Import: [ inet.0 ]

```

show route export instance detail

```

user@host> show route export instance detail
Instance: master                      Type: forwarding
  Flags: <config auto-policy> Options: <unicast multicast>
  Import policy: [ (ospf-master-from-red || isis-master-from-black) ]
Instance: black                       Type: non-forwarding
Instance: red                         Type: non-forwarding

```


show route export vrf-target

Syntax	show route export vrf-target <brief detail> <community <i>community--regular-expression</i> > <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display the VPN routing and forwarding (VRF) target communities for which policy-based route export is currently distributing routes. This command is relevant when there are overlapping virtual private networks (VPNs).
Options	<p>none—Display standard information about all target communities.</p> <p>brief detail—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>community <i>community-regular-expression</i>—(Optional) Display information about the specified community.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show route export vrf-target on page 830 show route export vrf-target community on page 830 show route export vrf-target detail on page 830
Output Fields	Table 131 on page 829 lists the output fields for the show route export vrf-target command. Output fields are listed in the approximate order in which they appear.

Table 131: show route export vrf-target Output Fields

Field Name	Field Description	Level of Output
Route target	Target communities for which auto-export is currently distributing routes.	brief none
Family	Routing table entries for the specified family.	brief none
<i>type-of-routing-table(s)</i>	Type of routing tables the feature handles: <ul style="list-style-type: none"> unicast—Indicates <i>instance.inet.0</i>. multicast—Indicates <i>instance.inet.2</i>. unicast multicast—Indicates <i>instance.inet.0</i> and <i>instance.inet.2</i>. 	brief none
Import	Number of routing tables that are currently importing routes with this target community. Omitted for tables that are not importing routes.	brief none

Table 131: show route export vrf-target Output Fields (*continued*)

Field Name	Field Description	Level of Output
Export	Number of routing tables that are currently exporting routes with this target community. Omitted for tables that are not exporting routes.	brief none
Target	Target communities, family, and options for which auto-export is currently distributing routes.	detail
Import table(s)	Name of the routing tables that are importing a particular route target.	detail
Export table(s)	Name of the routing tables that are exporting a particular route target.	detail

Sample Output

show route export vrf-target

```

user@host> show route export vrf-target
Route Target          Family      Import      Export
69:1                  inet        unicast      2           2
69:2                  inet        unicast      2           2

```

show route export vrf-target community

```

user@host> show route export vrf-target community target:69:1
Route Target          Family      Import      Export
69:1                  inet        unicast      2           2

```

show route export vrf-target detail

```

user@host> show route export vrf-target detail
Target: 1:12          inet        unicast
  Import table(s): vrf-11.inet.0 vrf-12.inet.0
  Export table(s): vrf-12.inet.0
Target: 1:13          inet        unicast
  Import table(s): vrf-12.inet.0 vrf-13.inet.0
  Export table(s): vrf-13.inet.0

```


show route extensive

Syntax	show route extensive <destination-prefix> <logical-system (all logical-system-name)>
Syntax (EX Series Switches)	show route extensive <destination-prefix>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display extensive information about the active entries in the routing tables.
Options	<p>none—Display all active entries in the routing table.</p> <p>destination-prefix—(Optional) Display active entries for the specified address or range of addresses.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show route extensive on page 838 show route extensive (Access Route) on page 844 show route extensive (BGP PIC Edge) on page 845 show route extensive (FRR and LFA) on page 845 show route extensive (Route Reflector) on page 846 show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 846 show route label detail (Multipoint LDP with Multicast-Only Fast Reroute) on page 847
Output Fields	Table 132 on page 831 describes the output fields for the show route extensive command. Output fields are listed in the approximate order in which they appear.

Table 132: show route extensive Output Fields

Field Name	Field Description
<i>routing-table-name</i>	Name of the routing table (for example, inet.0).
<i>number destinations</i>	Number of destinations for which there are routes in the routing table.
<i>number routes</i>	Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> active (routes that are active). holddown (routes that are in the pending state before being declared inactive). hidden (routes that are not used because of a routing policy).

Table 132: show route extensive Output Fields (*continued*)

Field Name	Field Description
route-destination (entry, announced)	<p>Route destination (for example: 10.0.0.1/24). The entry value is the number of route for this destination, and the announced value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:</p> <ul style="list-style-type: none"> • MPLS-label (for example, 80001). • interface-name (for example, ge-1/0/2). • neighbor-address:control-word-status:encapsulation type:vc-id:source (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96). • neighbor-address—Address of the neighbor. • control-word-status—Whether the use of the control word has been negotiated for this virtual circuit: NoCtrlWord or CtrlWord. • encapsulation type—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport. • vc-id—Virtual circuit identifier. • source—Source of the advertisement: Local or Remote.
TSI	Protocol header information.
label stacking	<p>(Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels.</p> <ul style="list-style-type: none"> • S=0 route indicates that a packet with an incoming label stack depth of two or more exits this router with one fewer label (the label-popping operation is performed). • If there is no S= information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).
[protocol, preference]	<p>Protocol from which the route was learned and the preference value for the route.</p> <ul style="list-style-type: none"> • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • -—A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route. <p>In every routing metric except for the BGP LocalPref attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the LocalPref value in the Preference2 field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the LocalPref value for Route 2 is 155, the Preference2 value is -156. Route 2 is preferred because it has a higher LocalPref value and a lower Preference2 value.</p>
Level	<p>(IS-IS only). In IS-IS, a single autonomous system (AS) can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.</p>
Route Distinguisher	IP subnet augmented with a 64-bit prefix.

Table 132: show route extensive Output Fields (*continued*)

Field Name	Field Description
PMSI	Provider multicast service interface (MVPN routing table).
Next-hop type	Type of next hop. For a description of possible values for this field, see the Output Field table in the show route detail command.
Next-hop reference count	Number of references made to the next hop.
Flood nexthop branches exceed maximum message	Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.
Source	IP address of the route source.
Next hop	Network layer address of the directly reachable neighboring system.
via	<p>Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word Selected. This field can also contain the following information:</p> <ul style="list-style-type: none"> • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.
Label-switched-path lsp-path-name	Name of the LSP used to reach the next hop.
Label operation	MPLS label and operation occurring at this routing device. The operation can be pop (where a label is removed from the top of the stack), push (where another label is added to the label stack), or swap (where a label is replaced by another label).
Offset	Whether the metric has been increased or decreased by an offset value.
Interface	(Local only) Local interface name.
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to recursively derive a forwarding next hop.
label-operation	MPLS label and operation occurring at this routing device. The operation can be pop (where a label is removed from the top of the stack), push (where another label is added to the label stack), or swap (where a label is replaced by another label).

Table 132: show route extensive Output Fields (*continued*)

Field Name	Field Description
Indirect next hops	<p>When present, a list of nodes that are used to resolve the path to the next-hop destination, in the order that they are resolved.</p> <p>When BGP PIC Edge is enabled, the output lines that contain Indirect next hop: weight follow next hops that the software can use to repair paths where a link failure occurs. The next-hop weight has one of the following values:</p> <ul style="list-style-type: none">• 0x1 indicates active next hops.• 0x4000 indicates passive next hops.
State	State of the route (a route can be in more than one state). See the Output Field table in the show route detail command.
Session ID	The BFD session ID number that represents the protection using MPLS fast reroute (FRR) and loop-free alternate (LFA).
Weight	<p>Weight for the backup path. If the weight of an indirect next hop is larger than zero, the weight value is shown.</p> <p>For sample output, see show route table.</p>

Table 132: show route extensive Output Fields (*continued*)

Field Name	Field Description
Inactive reason	<p>If the route is inactive, the reason for its current state is indicated. Typical reasons include:</p> <ul style="list-style-type: none"> • Active preferred—Currently active route was selected over this route. • Always compare MED—Path with a lower multiple exit discriminator (MED) is available. • AS path—Shorter AS path is available. • Cisco Non-deterministic MED selection—Cisco nondeterministic MED is enabled and a path with a lower MED is available. • Cluster list length—Path with a shorter cluster list length is available. • Forwarding use only—Path is only available for forwarding purposes. • IGP metric—Path through the next hop with a lower IGP metric is available. • IGP metric type—Path with a lower OSPF link-state advertisement type is available. • Interior > Exterior > Exterior via Interior—Direct, static, IGP, or EBGp path is available. • Local preference—Path with a higher local preference value is available. • Next hop address—Path with a lower metric next hop is available. • No difference—Path from a neighbor with a lower IP address is available. • Not Best in its group—Occurs when multiple peers of the same external AS advertise the same prefix and are grouped together in the selection process. When this reason is displayed, an additional reason is provided (typically one of the other reasons listed). • Number of gateways—Path with a higher number of next hops is available. • Origin—Path with a lower origin code is available. • OSPF version—Path does not support the indicated OSPF version. • RIB preference—Route from a higher-numbered routing table is available. • Route distinguisher—64-bit prefix added to IP subnets to make them unique. • Route metric or MED comparison—Route with a lower metric or MED is available. • Route preference—Route with a lower preference value is available. • Router ID—Path through a neighbor with a lower ID is available. • Unusable path—Path is not usable because of one of the following conditions: the route is damped, the route is rejected by an import policy, or the route is unresolved. • Update source—Last tiebreaker is the lowest IP address value.
Local AS	Autonomous system (AS) number of the local routing device.
Age	How long the route has been known.
AI GP	Accumulated interior gateway protocol (AIGP) BGP attribute.
Metric	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
MED-plus-IGP	Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.
TTL-Action	<p>For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances.</p> <p>For sample output, see show route table.</p>

Table 132: show route extensive Output Fields (*continued*)

Field Name	Field Description
Task	Name of the protocol that has added the route.
Announcement bits	List of protocols that announce this route. n-Resolve inet indicates that the route is used for route resolution for next hops found in the routing table. n is an index used by Juniper Networks customer support only.
AS path	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • Recorded—The AS path is recorded by the sample process (sampled). • ?—Incomplete; typically, the AS path was aggregated. <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> • []—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured. • { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • ()—Parentheses enclose a confederation. • ([])—Parentheses and brackets enclose a confederation set. <p>NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p>
validation-state	<p>(BGP-learned routes) Validation status of the route:</p> <ul style="list-style-type: none"> • Invalid—Indicates that the prefix is found, but either the corresponding AS received from the EBGp peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • Unknown—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • Unverified—Indicates that origin validation is not enabled for the BGP peers. • Valid—Indicates that the prefix and autonomous system pair are found in the database.
FECs bound to route	Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.
AS path: I <Originator>	(For route reflected output only) Originator ID attribute set by the route reflector.
Primary Upstream	When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path.
RPF Nexthops	When multipoint LDP with MoFRR is configured, the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.

Table 132: show route extensive Output Fields (*continued*)

Field Name	Field Description
Label	Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.
weight	Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.
VC Label	MPLS label assigned to the Layer 2 circuit virtual connection.
MTU	Maximum transmission unit (MTU) of the Layer 2 circuit.
VLAN ID	VLAN identifier of the Layer 2 circuit.
Cluster list	(For route reflected output only) Cluster ID sent by the route reflector.
Originator ID	(For route reflected output only) Address of router that originally sent the route to the route reflector.
Prefixes bound to route	Forwarding equivalent class (FEC) bound to this route. Applicable only to routes installed by LDP.
Communities	Community path attribute for the route. See the Output Field table in the show route detail command for all possible values for this field.
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).
control flags	Control flags: none or Site Down.
mtu	Maximum transmission unit (MTU) information.
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.
status vector	Layer 2 VPN and VPLS network layer reachability information (NLRI).
Localpref	Local preference value included in the route.
Router ID	BGP router ID as advertised by the neighbor in the open message.
Primary Routing Table	In a routing table group, the name of the primary routing table in which the route resides.
Secondary Tables	In a routing table group, the name of one or more secondary tables in which the route resides.
Originating RIB	Name of the routing table whose active route was used to determine the forwarding next-hop entry in the resolution database. For example, in the case of inet.0 resolving through inet.0 and inet.3, this field indicates which routing table, inet.0 or inet.3, provided the best path for a particular prefix.
Node path count	Number of nodes in the path.

Table 132: show route extensive Output Fields (*continued*)

Field Name	Field Description
Forwarding nexthops	Number of forwarding next hops. The forwarding next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.

Sample Output

show route extensive

```

user@host> show route extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
    *Static Preference: 5
        Next-hop reference count: 29
        Next hop: 192.168.71.254 via fxp0.0, selected
        State: <Active NoReadvrt Int Ext>
        Local AS: 69
        Age: 1:34:06
        Task: RT
        Announcement bits (2): 0-KRT 3-Resolve tree 2
        AS path: I

10.31.1.0/30 (2 entries, 1 announced)
    *Direct Preference: 0
        Next hop type: Interface
        Next-hop reference count: 2
        Next hop: via so-0/3/0.0, selected
        State: <Active Int>
        Local AS: 69
        Age: 1:32:40
        Task: IF
        Announcement bits (1): 3-Resolve tree 2
        AS path: I
    OSPF Preference: 10
        Next-hop reference count: 1
        Next hop: via so-0/3/0.0, selected
        State: <Int>
        Inactive reason: Route Preference
        Local AS: 69
        Age: 1:32:40 Metric: 1
        Area: 0.0.0.0
        Task: OSPF
        AS path: I

10.31.1.1/32 (1 entry, 1 announced)
    *Local Preference: 0
        Next hop type: Local
        Next-hop reference count: 7
        Interface: so-0/3/0.0
        State: <Active NoReadvrt Int>
        Local AS: 69
        Age: 1:32:43
        Task: IF
        Announcement bits (1): 3-Resolve tree 2
        AS path: I

```



```

...

10.31.2.0/30 (1 entry, 1 announced)
TSI:
KRT in-kerne1 10.31.2.0/30 -> {10.31.1.6}
    *OSPF   Preference: 10
            Next-hop reference count: 9
            Next hop: via so-0/3/0.0
            Next hop: 10.31.1.6 via ge-3/1/0.0, selected
            State: <Active Int>
            Local AS:    69
            Age: 1:32:19   Metric: 2
            Area: 0.0.0.0
            Task: OSPF
            Announcement bits (2): 0-KRT 3-Resolve tree 2
            AS path: I

...

224.0.0.2/32 (1 entry, 1 announced)
TSI:
KRT in-kerne1 224.0.0.2/32 -> {}
    *PIM    Preference: 0
            Next-hop reference count: 18
            State: <Active NoReadvrt Int>
            Local AS:    69
            Age: 1:34:08
            Task: PIM Recv
            Announcement bits (2): 0-KRT 3-Resolve tree 2
            AS path: I

...

224.0.0.22/32 (1 entry, 1 announced)
TSI:
KRT in-kerne1 224.0.0.22/32 -> {}
    *IGMP   Preference: 0
            Next-hop reference count: 18
            State: <Active NoReadvrt Int>
            Local AS:    69
            Age: 1:34:06
            Task: IGMP
            Announcement bits (2): 0-KRT 3-Resolve tree 2
            AS path: I

...

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.255.70.103/32 (1 entry, 1 announced)
State: <FlashAll>
    *RSVP   Preference: 7
            Next-hop reference count: 6
            Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
            Label-switched-path green-r1-r3
            Label operation: Push 100096
            State: <Active Int>
            Local AS:    69
            Age: 1:28:12   Metric: 2
            Task: RSVP
            Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
            AS path: I

```

```

10.255.71.238/32 (1 entry, 1 announced)
  State: <FlashAll>
  *RSVP   Preference: 7
          Next-hop reference count: 6
          Next hop: via so-0/3/0.0 weight 0x1, selected
          Label-switched-path green-r1-r2
          State: <Active Int>
          Local AS: 69
          Age: 1:28:12    Metric: 1
          Task: RSVP
          Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
          AS path: I

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

...

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
  *Direct Preference: 0
          Next hop type: Interface
          Next-hop reference count: 1
          Next hop: via lo0.0, selected
          State: <Active Int>
          Local AS: 69
          Age: 1:34:07
          Task: IF
          AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

0 (1 entry, 1 announced)
TSI:
KRT in-kernel 0 /36 -> {}
  *MPLS   Preference: 0
          Next hop type: Receive
          Next-hop reference count: 6
          State: <Active Int>
          Local AS: 69
          Age: 1:34:08    Metric: 1
          Task: MPLS
          Announcement bits (1): 0-KRT
          AS path: I

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
299840 (1 entry, 1 announced)
TSI:
KRT in-kernel 299840 /52 -> {indirect(1048575)}
  *RSVP   Preference: 7/2
          Next hop type: Flood
          Address: 0x9174a30
          Next-hop reference count: 4
          Next hop type: Router, Next hop index: 798
          Address: 0x9174c28
          Next-hop reference count: 2
          Next hop: 8.0.0.2 via lt-1/2/0.9 weight 0x1
          Label-switched-path R2-to-R4-2p2mp

```

```

Label operation: Pop
Next hop type: Router, Next hop index: 1048574
Address: 0x92544f0
Next-hop reference count: 2
Next hop: 7.0.0.2 via lt-1/2/0.7 weight 0x1
Label-switched-path R2-to-R200-p2mp
Label operation: Pop
Next hop: 6.0.0.2 via lt-1/2/0.5 weight 0x8001
Label operation: Pop
State: <Active Int>
Age: 1:29      Metric: 1
Task: RSVP
Announcement bits (1): 0-KRT
AS path: I...

```

800010 (1 entry, 1 announced)

TSI:

```

KRT in-kernel 800010 /36 -> {vt-3/2/0.32769}
  *VPLS Preference: 7
    Next-hop reference count: 2
    Next hop: via vt-3/2/0.32769, selected
    Label operation: Pop
    State: <Active Int>
    Age: 1:31:53
    Task: Common L2 VC
    Announcement bits (1): 0-KRT
    AS path: I

```

vt-3/2/0.32769 (1 entry, 1 announced)

TSI:

```

KRT in-kernel vt-3/2/0.32769.0 /16 -> {indirect(1048574)}
  *VPLS Preference: 7
    Next-hop reference count: 2
    Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
    Label-switched-path green-r1-r3
    Label operation: Push 800012, Push 100096(top)
    Protocol next hop: 10.255.70.103
    Push 800012
    Indirect next hop: 87272e4 1048574
    State: <Active Int>
    Age: 1:31:53      Metric2: 2
    Task: Common L2 VC
    Announcement bits (2): 0-KRT 1-Common L2 VC
    AS path: I
    Communities: target:11111:1 Layer2-info: encaps:VPLS,
    control flags:, mtu: 0
    Indirect next hops: 1
      Protocol next hop: 10.255.70.103 Metric: 2
      Push 800012
      Indirect next hop: 87272e4 1048574
      Indirect path forwarding next hops: 1
        Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1
        10.255.70.103/32 Originating RIB: inet.3
        Metric: 2      Node path count: 1
        Forwarding nexthops: 1
        Nexthop: 10.31.1.6 via ge-3/1/0.0

```

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

abcd::10:255:71:52/128 (1 entry, 0 announced)

```
*Direct Preference: 0
  Next hop type: Interface
  Next-hop reference count: 1
  Next hop: via lo0.0, selected
  State: <Active Int>
  Local AS: 69
  Age: 1:34:07
  Task: IF
  AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:34:07
    Task: IF
    AS path: I

ff02::2/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::2/128 -> {}
  *PIM Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:34:08
    Task: PIM Recv6
    Announcement bits (1): 0-KRT
    AS path: I

ff02::d/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::d/128 -> {}
  *PIM Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:34:08
    Task: PIM Recv6
    Announcement bits (1): 0-KRT
    AS path: I

ff02::16/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::16/128 -> {}
  *MLD Preference: 0
    Next-hop reference count: 18
    State: <Active NoReadvrt Int>
    Local AS: 69
    Age: 1:34:06
    Task: MLD
    Announcement bits (1): 0-KRT
    AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
  *Direct Preference: 0
```

```

Next hop type: Interface
Next-hop reference count: 1
Next hop: via lo0.16385, selected
State: <Active NoReadvrt Int>
Age: 1:34:07
Task: IF
AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

10.255.70.103:1:3:1/96 (1 entry, 1 announced)
  *BGP Preference: 170/-101
    Route Distinguisher: 10.255.70.103:1
    Next-hop reference count: 7
    Source: 10.255.70.103
    Protocol next hop: 10.255.70.103
    Indirect next hop: 2 no-forward
    State: <Secondary Active Int Ext>
    Local AS: 69 Peer AS: 69
    Age: 1:28:12 Metric2: 1
    Task: BGP_69.10.255.70.103+179
    Announcement bits (1): 0-green-l2vpn
    AS path: I
    Communities: target:11111:1 Layer2-info: encaps:VPLS,
    control flags:, mtu: 0
    Label-base: 800008, range: 8
    Localpref: 100
    Router ID: 10.255.70.103
    Primary Routing Table bgp.l2vpn.0

10.255.71.52:1:1:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699540
  *L2VPN Preference: 170/-1
    Next-hop reference count: 5
    Protocol next hop: 10.255.71.52
    Indirect next hop: 0 -
    State: <Active Int Ext>
    Age: 1:34:03 Metric2: 1
    Task: green-l2vpn
    Announcement bits (1): 1-BGP.0.0.0.0+179
    AS path: I
    Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
    mtu: 0
    Label-base: 800016, range: 8, status-vector: 0x9F

10.255.71.52:1:5:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699528
  *L2VPN Preference: 170/-101
    Next-hop reference count: 5
    Protocol next hop: 10.255.71.52
    Indirect next hop: 0 -
    State: <Active Int Ext>
    Age: 1:34:03 Metric2: 1
    Task: green-l2vpn
    Announcement bits (1): 1-BGP.0.0.0.0+179
    AS path: I
    Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
    Label-base: 800008, range: 8, status-vector: 0x9F

```

```
...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

TSI:

10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
  *L2CKT Preference: 7
    Next hop: via so-1/1/2.0 weight 1, selected
    Label-switched-path my-lsp
    Label operation: Push 100000[0]
    Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
    State: <Active Int>
    Local AS: 99
    Age: 10:21
    Task: l2 circuit
    Announcement bits (1): 0-LDP
    AS path: I
    VC Label 100000, MTU 1500, VLAN ID 512

55.0.0.0/24 (1 entry, 1 announced)
TSI:
KRT queued (pending) add
  55.0.0.0/24 -> {Push 300112}
    *BGP Preference: 170/-101
      Next hop type: Router
      Address: 0x925c208
      Next-hop reference count: 2
      Source: 10.0.0.9
      Next hop: 10.0.0.9 via ge-1/2/0.15, selected
      Label operation: Push 300112
      Label TTL action: prop-ttl
      State: <Active Ext>
      Local AS: 7019 Peer AS: 13979
      Age: 1w0d 23:06:56
      AIGP: 25
      Task: BGP_13979.10.0.0.9+56732
      Announcement bits (1): 0-KRT
      AS path: 13979 7018 I
      Accepted
      Route Label: 300112
      Localpref: 100
      Router ID: 10.9.9.1
```

show route extensive (Access Route)

```
user@host> show route 13.160.0.102 extensive
inet.0: 39256 destinations, 39258 routes (39255 active, 0 holddown, 1 hidden)
13.160.0.102/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 13.160.0.102/32 -> {13.160.0.2}
OSPF area : 0.0.0.0, LSA ID : 13.160.0.102, LSA type : Extern
  *Access Preference: 13
    Next-hop reference count: 78472
    Next hop: 13.160.0.2 via fe-0/0/0.0, selected
    State: <Active Int>
  Age: 12
    Task: RPD Unix Domain Server./var/run/rpd_serv.local
    Announcement bits (2): 0-KRT 1-OSPFv2
    AS path: I
```

show route extensive (BGP PIC Edge)

```

user@host> show route 1.1.1.6 extensive
ed.inet.0: 6 destinations, 9 routes (6 active, 0 holddown, 0 hidden)
  1.1.1.6/32 (3 entries, 2 announced)
    State: <CalcForwarding>
    TSI:
    KRT in-kerne1 1.1.1.6/32 -> {indirect(1048574), indirect(1048577)}
    Page 0 idx 0 Type 1 val 9219e30
    Nexthop: Self
    AS path: [2] 3 I
    Communities: target:2:1
    Path 1.1.1.6 from 1.1.1.4 Vector len 4. Val: 0
  ..
    #Multipath Preference: 255
    Next hop type: Indirect
    Address: 0x93f4010
    Next-hop reference count: 2
  ..
    Protocol next hop: 1.1.1.4
    Push 299824
    Indirect next hop: 944c000 1048574 INH Session ID: 0x3
    Indirect next hop: weight 0x1
    Protocol next hop: 1.1.1.5
    Push 299824
    Indirect next hop: 944c1d8 1048577 INH Session ID: 0x4
    Indirect next hop: weight 0x4000
    State: <ForwardingOnly Int Ext>
    Inactive reason: Forwarding use only
    Age: 25      Metric2: 15
    Validation State: unverified
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: 3 I
    Communities: target:2:1

```

show route extensive (FRR and LFA)

```

user@host> show route 20.31.2.0 extensive
inet.0: 46 destinations, 49 routes (45 active, 0 holddown, 1 hidden)
  20.31.2.0/24 (2 entries, 1 announced)
    State: FlashAll
    TSI:
    KRT in-kerne1 20.31.2.0/24 -> {Push 299776, Push 299792}
    *RSVP Preference: 7/1
    Next hop type: Router, Next hop index: 1048574
    Address: 0xbbbc010
    Next-hop reference count: 5
    Next hop: 10.31.1.2 via ge-2/1/8.0 weight 0x1, selected
    Label-switched-path europa-d-to-europa-e
    Label operation: Push 299776
    Label TTL action: prop-ttl
    Session Id: 0x201
    Next hop: 10.31.2.2 via ge-2/1/4.0 weight 0x4001
    Label-switched-path europa-d-to-europa-e
    Label operation: Push 299792
    Label TTL action: prop-ttl
    Session Id: 0x202
    State: Active Int
    Local AS: 100
    Age: 5:31 Metric: 2

```

```

Task: RSVP
Announcement bits (1): 0-KRT
AS path: I
OSPF Preference: 10
Next hop type: Router, Next hop index: 615
Address: 0xb9d78c4
Next-hop reference count: 7
Next hop: 10.31.1.2 via ge-2/1/8.0, selected
Session Id: 0x201
State: Int
Inactive reason: Route Preference
Local AS: 100
Age: 5:35 Metric: 3
Area: 0.0.0.0
Task: OSPF
AS path: I

```

show route extensive (Route Reflector)

```

user@host> show route extensive
1.0.0.0/8 (1 entry, 1 announced)

TSI:
KRT in-kernel 1.0.0.0/8 -> {indirect(40)}
*BGP Preference: 170/-101
Source: 192.168.4.214
Protocol next hop: 207.17.136.192 Indirect next hop: 84ac908 40
State: <Active Int Ext>
Local AS: 10458 Peer AS: 10458
Age: 3:09 Metric: 0 Metric2: 0
Task: BGP_10458.192.168.4.214+1033
Announcement bits (2): 0-KRT 4-Resolve inet.0
AS path: 3944 7777 I <Originator>
Cluster list: 1.1.1.1
Originator ID: 10.255.245.88
Communities: 7777:7777
Localpref: 100
Router ID: 4.4.4.4
Indirect next hops: 1
    Protocol next hop: 207.17.136.192 Metric: 0
    Indirect next hop: 84ac908 40
    Indirect path forwarding next hops: 0
    Next hop type: Discard

```

show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```

user@host> show route label 299872 detail
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
*LDP Preference: 9
Next hop type: Flood
Next-hop reference count: 3
Address: 0x9097d90
Next hop: via vt-0/1/0.1
Next-hop index: 661
Label operation: Pop
Address: 0x9172130
Next hop: via so-0/0/3.0
Next-hop index: 654
Label operation: Swap 299872
State: **Active Int>

```



```

Local AS: 1001
Age: 8:20      Metric: 1
Task: LDP
Announcement bits (1): 0-KRT
AS path: I
FECs bound to route: P2MP root-addr 10.255.72.166, grp 232.1.1.1,
src 192.168.142.2

```

show route label detail (Multipoint LDP with Multicast-Only Fast Reroute)

```
user@host> show route label 301568 detail
```

```

mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
  *LDP    Preference: 9
    Next hop type: Flood
    Address: 0x2735208
    Next-hop reference count: 3
    Next hop type: Router, Next hop index: 1397
    Address: 0x2735d2c
    Next-hop reference count: 3
    Next hop: 1.3.8.2 via ge-1/2/22.0
    Label operation: Pop
    Load balance label: None;
    Next hop type: Router, Next hop index: 1395
    Address: 0x2736290
    Next-hop reference count: 3
    Next hop: 1.3.4.2 via ge-1/2/18.0
    Label operation: Pop
    Load balance label: None;
    State: <Active Int AckRequest MulticastRPF>
    Local AS: 10
    Age: 54:05      Metric: 1
    Validation State: unverified
    Task: LDP
    Announcement bits (1): 0-KRT
    AS path: I
    FECs bound to route: P2MP root-addr 1.1.1.1, grp: 232.1.1.1, src:
192.168.219.11
      Primary Upstream : 1.1.1.3:0--1.1.1.2:0
        RPF Nexthops :
          ge-1/2/15.0, 1.2.94.1, Label: 301568, weight: 0x1
          ge-1/2/14.0, 1.2.3.1, Label: 301568, weight: 0x1
      Backup Upstream : 1.1.1.3:0--1.1.1.6:0
        RPF Nexthops :
          ge-1/2/20.0, 1.2.96.1, Label: 301584, weight: 0xffffe
          ge-1/2/19.0, 1.3.6.1, Label: 301584, weight: 0xffffe

```

show route flow validation

Syntax	show route flow validation <brief detail> <ip-prefix> <table table-name> <logical-system (all logical-system-name)>
Syntax (EX Series Switches)	show route flow validation <brief detail> <ip-prefix> <table table-name>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display flow route information.
Options	<p>none—Display flow route information.</p> <p>brief detail—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>ip-prefix—(Optional) IP address for the flow route.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>table table-name—(Optional) Display flow route information for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the show route flow validation inet command).</p>
Required Privilege Level	view
List of Sample Output	show route flow validation on page 849
Output Fields	Table 133 on page 848 lists the output fields for the show route flow validation command. Output fields are listed in the approximate order in which they appear.

Table 133: show route flow validation Output Fields

Field Name	Field Description	Level of Output
<i>routing-table-name</i>	Name of the routing table (for example, inet.0).	All levels
<i>prefix</i>	Route address.	All levels
Active unicast route	Active route in the routing table.	All levels
Dependent flow destinations	Number of flows for which there are routes in the routing table.	All levels

Table 133: show route flow validation Output Fields (*continued*)

Field Name	Field Description	Level of Output
Origin	Source of the route flow.	All levels
Neighbor AS	Autonomous system identifier of the neighbor.	All levels
Flow destination	Number of entries and number of destinations that match the route flow.	All levels
Unicast best match	Destination that is the best match for the route flow.	All levels
Flags	Information about the route flow.	All levels

Sample Output

show route flow validation

```
user@host> show route flow validation
inet.0:
10.0.5.0/24Active unicast route
Dependent flow destinations: 1
Origin: 192.168.224.218, Neighbor AS: 65001
Flow destination (3 entries, 1 match origin)
Unicast best match: 10.0.5.0/24
Flags: SubtreeApex Consistent
```

show route forwarding-table

Syntax	<pre>show route forwarding-table <detail extensive summary> <all> <ccc interface-name> <destination destination-prefix> <family family matching matching> <interface-name interface-name> <label name> <matching matching> <multicast> <table (default logical-system-name/routing-instance-name routing-instance-name)> <vlan (all vlan-name)> <vpn vpn></pre>
Syntax (MX Series Routers)	<pre>show route forwarding-table <detail extensive summary> <all> <bridge-domain (all domain-name)> <ccc interface-name> <destination destination-prefix> <family family matching matching> <interface-name interface-name> <label name> <learning-vlan-id learning-vlan-id> <matching matching> <multicast> <table (default logical-system-name/routing-instance-name routing-instance-name)> <vlan (all vlan-name)> <vpn vpn></pre>
Syntax (TX Matrix and TX Matrix Plus Routers)	<pre>show route forwarding-table <detail extensive summary> <all> <ccc interface-name> <destination destination-prefix> <family family matching matching> <interface-name interface-name> <matching matching> <label name> <lcc number> <multicast> <table routing-instance-name> <vpn vpn></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Option bridge-domain introduced in Junos OS Release 7.5</p> <p>Option learning-vlan-id introduced in Junos OS Release 8.4</p> <p>Options all and vlan introduced in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>

Description Display the Routing Engine's forwarding table, including the network-layer prefixes and their next hops. This command is used to help verify that the routing protocol process has relayed the correction information to the forwarding table. The Routing Engine constructs and maintains one or more routing tables. From the routing tables, the Routing Engine derives a table of active routes, called the forwarding table.



NOTE: The Routing Engine copies the forwarding table to the Packet Forwarding Engine, the part of the router that is responsible for forwarding packets. To display the entries in the Packet Forwarding Engine's forwarding table, use the **show pfe route** command.

Options **none**—Display the routes in the forwarding tables. By default, the **show route forwarding-table** command does not display information about private, or internal, forwarding tables.

detail | extensive | summary—(Optional) Display the specified level of output.

all—(Optional) Display routing table entries for all forwarding tables, including private, or internal, tables.

bridge-domain (all | *bridge-domain-name*)—(MX Series routers only) (Optional) Display route entries for all bridge domains or the specified bridge domain.

ccc *interface-name*—(Optional) Display route entries for the specified circuit cross-connect interface.

destination *destination-prefix*—(Optional) Destination prefix.

family *family*—(Optional) Display routing table entries for the specified family: **fibre-channel**, **fmembers**, **inet**, **inet6**, **iso**, **mpls**, **tnp**, **unix**, **vpls**, or **vlan-classification**.

interface-name *interface-name*—(Optional) Display routing table entries for the specified interface.

label *name*—(Optional) Display route entries for the specified label.

lcc *number*—(TX Matrix and TX matrix Plus routers only) (Optional) On a routing matrix composed of a TX Matrix router and T640 routers, display information for the specified T640 router (or line-card chassis) connected to the TX Matrix router. On a routing matrix composed of the TX Matrix Plus router and T1600 or T4000 routers, display information for the specified router (line-card chassis) connected to the TX Matrix Plus router.

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

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.

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

learning-vlan-id *learning-vlan-id*—(MX Series routers only) (Optional) Display learned information for all VLANs or for the specified VLAN.

matching *matching*—(Optional) Display routing table entries matching the specified prefix or prefix length.

multicast—(Optional) Display routing table entries for multicast routes.

table (default | *logical-system-name/routing-instance-name* | *routing-instance-name*)—(Optional) Display route entries for all the routing tables in the main routing instance or for the specified routing instance. If your device supports logical systems, you can also display route entries for the specified logical system and routing instance. To view the routing instances on your device, use the [show route instance](#) command.

vlan (all | *vlan-name*)—(Optional) Display information for all VLANs or for the specified VLAN.

vpn *vpn*—(Optional) Display routing table entries for a specified VPN.

Required Privilege Level

view

List of Sample Output

[show route forwarding-table on page 855](#)
[show route forwarding-table detail on page 856](#)
[show route forwarding-table destination extensive \(Weights and Balances\) on page 856](#)
[show route forwarding-table extensive on page 857](#)
[show route forwarding-table extensive \(RPF\) on page 858](#)
[show route forwarding-table family mpls on page 859](#)
[show route forwarding-table family vpls on page 859](#)
[show route forwarding-table vpls \(Broadcast, unknown unicast, and multicast \(BUM\) hashing is enabled\) on page 859](#)
[show route forwarding-table vpls \(Broadcast, unknown unicast, and multicast \(BUM\) hashing is enabled with MAC Statistics\) on page 860](#)
[show route forwarding-table family vpls extensive on page 860](#)
[show route forwarding-table table default on page 861](#)
[show route forwarding-table table logical-system-name/routing-instance-name on page 862](#)
[show route forwarding-table vpn on page 863](#)

Output Fields

[Table 134 on page 853](#) lists the output fields for the **show route forwarding-table** command. Output fields are listed in the approximate order in which they appear. Field names might be abbreviated (as shown in parentheses) when no level of output is specified, or when the **detail** keyword is used instead of the **extensive** keyword.

Table 134: show route forwarding-table Output Fields

Field Name	Field Description	Level of Output
Logical system	Name of the logical system. This field is displayed if you specify the table <i>logical-system-name/routing-instance-name</i> option on a device that is configured for and supports logical systems.	All levels
Routing table	Name of the routing table (for example, inet, inet6, mpls).	All levels
Address family	Address family (for example, IP, IPv6, ISO, MPLS, and VPLS).	All levels
Destination	Destination of the route.	detail extensive
Route Type (Type)	How the route was placed into the forwarding table. When the detail keyword is used, the route type might be abbreviated (as shown in parentheses): <ul style="list-style-type: none"> • cloned (clon)—(TCP or multicast only) Cloned route. • destination (dest)—Remote addresses directly reachable through an interface. • destination down (iddn)—Destination route for which the interface is unreachable. • interface cloned (ifcl)—Cloned route for which the interface is unreachable. • route down (ifdn)—Interface route for which the interface is unreachable. • ignore (ignr)—Ignore this route. • interface (intf)—Installed as a result of configuring an interface. • permanent (perm)—Routes installed by the kernel when the routing table is initialized. • user—Routes installed by the routing protocol process or as a result of the configuration. 	All levels
Route Reference (RtRef)	Number of routes to reference.	detail extensive
Flags	Route type flags: <ul style="list-style-type: none"> • none—No flags are enabled. • accounting—Route has accounting enabled. • cached—Cache route. • incoming-iface interface-number—Check against incoming interface. • prefix load balance—Load balancing is enabled for this prefix. • rt nh decoupled—Route has been decoupled from the next hop to the destination. • sent to PFE—Route has been sent to the Packet Forwarding Engine. • static—Static route. 	extensive
Next hop	IP address of the next hop to the destination.	detail extensive

Table 134: show route forwarding-table Output Fields (*continued*)

Field Name	Field Description	Level of Output
Next hop Type (Type)	<p>Next-hop type. When the detail keyword is used, the next-hop type might be abbreviated (as indicated in parentheses):</p> <ul style="list-style-type: none"> • broadcast (bcst)—Broadcast. • deny—Deny. • discard (dscd)—Discard. • hold—Next hop is waiting to be resolved into a unicast or multicast type. • indexed (idxd)—Indexed next hop. • indirect (indr)—Indirect next hop. • local (locl)—Local address on an interface. • routed multicast (mcrst)—Regular multicast next hop. • multicast (mcst)—Wire multicast next hop (limited to the LAN). • multicast discard (mdsc)—Multicast discard. • multicast group (mgrp)—Multicast group member. • receive (rcv)—Receive. • reject (rjct)—Discard. An ICMP unreachable message was sent. • resolve (rslv)—Resolving the next hop. • unicast (ucst)—Unicast. • unilist (ulst)—List of unicast next hops. A packet sent to this next hop goes to any next hop in the list. 	detail extensive
Index	Software index of the next hop that is used to route the traffic for a given prefix.	detail extensive none
Route interface-index	Logical interface index from which the route is learned. For example, for interface routes, this is the logical interface index of the route itself. For static routes, this field is zero. For routes learned through routing protocols, this is the logical interface index from which the route is learned.	extensive
Reference (NhRef)	Number of routes that refer to this next hop.	detail extensive none
Next-hop interface (Netif)	Interface used to reach the next hop.	detail extensive none
Weight	Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible (see the Balance field description).	extensive
Balance	Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a router is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.	extensive
RPF interface	List of interfaces from which the prefix can be accepted. Reverse path forwarding (RPF) information is displayed only when rpf-check is configured on the interface.	extensive

Sample Output

show route forwarding-table

```

user@host> show route forwarding-table
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                rjct   46    4
0.0.0.0/32       perm  0                dscd   44    1
1.1.1.0/24       ifdn  0                rslv   608   1 ge-2/0/1.0
1.1.1.0/32       iddn  0 1.1.1.0         recv   606   1 ge-2/0/1.0
1.1.1.1/32       user  0                rjct   46    4
1.1.1.1/32       intf  0 1.1.1.1         locl   607   2
1.1.1.1/32       iddn  0 1.1.1.1         locl   607   2
1.1.1.255/32     iddn  0 ff:ff:ff:ff:ff:ff bcst   605   1 ge-2/0/1.0
10.0.0.0/24      intf  0                rslv   616   1 ge-2/0/0.0
10.0.0.0/32      dest  0 10.0.0.0         recv   614   1 ge-2/0/0.0
10.0.0.1/32      intf  0 10.0.0.1         locl   615   2
10.0.0.1/32      dest  0 10.0.0.1         locl   615   2
10.0.0.255/32    dest  0 10.0.0.255       bcst   613   1 ge-2/0/0.0
10.1.1.0/24      ifdn  0                rslv   612   1 ge-2/0/1.0
10.1.1.0/32      iddn  0 10.1.1.0         recv   610   1 ge-2/0/1.0
10.1.1.1/32      user  0                rjct   46    4
10.1.1.1/32      intf  0 10.1.1.1         locl   611   2
10.1.1.1/32      iddn  0 10.1.1.1         locl   611   2
10.1.1.255/32    iddn  0 ff:ff:ff:ff:ff:ff bcst   609   1 ge-2/0/1.0
10.209.0.0/16    user  0 10.209.63.254    ucst   419   20 fxp0.0
10.209.0.0/16    user  1 0:12:1e:ca:98:0  ucst   419   20 fxp0.0
10.209.0.0/18    intf  0                rslv   418   1 fxp0.0
10.209.0.0/32    dest  0 10.209.0.0         recv   416   1 fxp0.0
10.209.2.131/32  intf  0 10.209.2.131     locl   417   2
10.209.2.131/32  dest  0 10.209.2.131     locl   417   2
10.209.17.55/32  dest  0 0:30:48:5b:78:d2 ucst   435   1 fxp0.0
10.209.63.42/32  dest  0 0:23:7d:58:92:ca ucst   434   1 fxp0.0
10.209.63.254/32 dest  0 0:12:1e:ca:98:0  ucst   419   20 fxp0.0
10.209.63.255/32 dest  0 10.209.63.255    bcst   415   1 fxp0.0
10.227.0.0/16    user  0 10.209.63.254    ucst   419   20 fxp0.0

...

Routing table: iso
ISO:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                rjct   27    1
47.0005.80ff.f800.0000.0108.0003.0102.5524.5220.00
intf  0                locl   28    1

Routing table: inet6
Internet6:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                rjct   6    1
ff00::/8         perm  0                mdsc   4    1
ff02::1/128      perm  0 ff02::1            mcst   3    1

Routing table: ccc
MPLS:
Interface.Label  Type RtRef Next hop          Type Index NhRef Netif
default          perm  0                rjct  16    1
100004(top)fe-0/0/1.0

```

show route forwarding-table detail

```

user@host> show route forwarding-table detail
Routing table: inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          user   2 0:90:69:8e:b1:1b ucst  132   4 fxp0.0
default          perm   0                               rjct   14    1
10.1.1.0/24      intf   0 ff.3.0.21         ucst  322   1 so-5/3/0.0
10.1.1.0/32      dest   0 10.1.1.0          recv  324   1 so-5/3/0.0
10.1.1.1/32      intf   0 10.1.1.1          locl  321   1
10.1.1.255/32    dest   0 10.1.1.255        bcst  323   1 so-5/3/0.0
10.21.21.0/24    intf   0 ff.3.0.21         ucst  326   1 so-5/3/0.0
10.21.21.0/32    dest   0 10.21.21.0        recv  328   1 so-5/3/0.0
10.21.21.1/32    intf   0 10.21.21.1        locl  325   1
10.21.21.255/32  dest   0 10.21.21.255      bcst  327   1 so-5/3/0.0
127.0.0.1/32     intf   0 127.0.0.1         locl  320   1
172.17.28.19/32  clon   1 192.168.4.254     ucst  132   4 fxp0.0
172.17.28.44/32  clon   1 192.168.4.254     ucst  132   4 fxp0.0

...

Routing table: private1__inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm   0                               rjct   46    1
10.0.0.0/8       intf   0                               rslv  136   1 fxp1.0
10.0.0.0/32      dest   0 10.0.0.0          recv  134   1 fxp1.0
10.0.0.4/32      intf   0 10.0.0.4          locl  135   2
10.0.0.4/32      dest   0 10.0.0.4          locl  135   2

...

Routing table: iso
ISO:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm   0                               rjct   38    1

Routing table: inet6
Internet6:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm   0                               rjct   22    1
ff00::/8         perm   0                               mdsc   21    1
ff02::1/128      perm   0 ff02::1          mcst   17    1

...

Routing table: mpls
MPLS:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm   0                               rjct  28    1

```

show route forwarding-table destination extensive (Weights and Balances)

```

user@host> show route forwarding-table destination 3.4.2.1 extensive
Routing table: inet [Index 0]
Internet:

Destination: 3.4.2.1/32
Route type: user
Route reference: 0                               Route interface-index: 0

```

```

Flags: sent to PFE
Next-hop type: unilist          Index: 262143  Reference: 1
Nexthop: 4.4.4.4
Next-hop type: unicast         Index: 335      Reference: 2
Next-hop interface: so-1/1/0.0 Weight: 22     Balance: 3
Nexthop: 145.12.1.2
Next-hop type: unicast         Index: 337     Reference: 2
Next-hop interface: so-0/1/2.0 Weight: 33     Balance: 33

```

show route forwarding-table extensive

```

user@host> show route forwarding-table extensive
Routing table: inet [Index 0]
Internet:

Destination: default
Route type: user
Route reference: 2                      Route interface-index: 0
Flags: sent to PFE
Nexthop: 0:90:69:8e:b1:1b
Next-hop type: unicast                 Index: 132      Reference: 4
Next-hop interface: fxp0.0

Destination: default
Route type: permanent
Route reference: 0                      Route interface-index: 0
Flags: none
Next-hop type: reject                 Index: 14       Reference: 1

Destination: 127.0.0.1/32
Route type: interface
Route reference: 0                      Route interface-index: 0
Flags: sent to PFE
Nexthop: 127.0.0.1
Next-hop type: local                   Index: 320      Reference: 1

...

Routing table: private1__inet [Index 1]
Internet:

Destination: default
Route type: permanent
Route reference: 0                      Route interface-index: 0
Flags: sent to PFE
Next-hop type: reject                 Index: 46       Reference: 1

Destination: 10.0.0.0/8
Route type: interface
Route reference: 0                      Route interface-index: 3
Flags: sent to PFE
Next-hop type: resolve                 Index: 136      Reference: 1
Next-hop interface: fxp1.0

...

Routing table: iso [Index 0]
ISO:

Destination: default
Route type: permanent

```

```

Route reference: 0
Flags: sent to PFE
Next-hop type: reject
Route interface-index: 0
Index: 38      Reference: 1

Routing table: inet6 [Index 0]
Internet6:

Destination: default
Route type: permanent
Route reference: 0
Flags: sent to PFE
Next-hop type: reject
Route interface-index: 0
Index: 22      Reference: 1

Destination: ff00::/8
Route type: permanent
Route reference: 0
Flags: sent to PFE
Next-hop type: multicast discard
Route interface-index: 0
Index: 21      Reference: 1

...

Routing table: private1__inet6 [Index 1]
Internet6:

Destination: default
Route type: permanent
Route reference: 0
Flags: sent to PFE
Next-hop type: reject
Route interface-index: 0
Index: 54      Reference: 1

Destination: fe80::2a0:a5ff:fe3d:375/128
Route type: interface
Route reference: 0
Flags: sent to PFE
Next-hop: fe80::2a0:a5ff:fe3d:375
Next-hop type: local
Route interface-index: 0
Index: 75      Reference: 1

...

```

show route forwarding-table extensive (RPF)

The next example is based on the following configuration, which enables an RPF check on all routes that are learned from this interface, including the interface route:

```

so-1/1/0 {
  unit 0 {
    family inet {
      rpf-check;
      address 15.95.1.2/30;
    }
  }
}

```

```

user@host> show route forwarding-table extensive
Routing table: inet [Index 0]
Internet:
...
...
Destination: 15.95.1.3/32
Route type: destination
Route reference: 0
Route interface-index: 67

```

```

Flags: sent to PFE
Nexthop: 15.95.1.3
Next-hop type: broadcast          Index: 328      Reference: 1
Next-hop interface: so-1/1/0.0
RPF interface: so-1/1/0.0

```

show route forwarding-table family mpls

```

user@host> show route forwarding-table family mpls
Routing table: mpls
MPLS:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm  0
0                user  0
1                user  0
2                user  0
100000           user  0 10.31.1.6      swap 100001      fe-1/1/0.0
800002           user  0                Pop          vt-0/3/0.32770

vt-0/3/0.32770 (VPLS)
                  user  0                indr  351      4
                  Push 800000, Push 100002(top)

so-0/0/0.0

```

show route forwarding-table family vpls

```

user@host> show route forwarding-table family vpls
Routing table: green.vpls
VPLS:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          dymn  0
default          perm  0
fe-0/1/0.0       dymn  0
00:90:69:0c:20:1f/48      <<<<<Remote CE

                  dymn  0                indr  351      4
                  Push 800000, Push 100002(top)

so-0/0/0.0
00:90:69:85:b0:1f/48      <<<<<Local CE

                  dymn  0                ucst  354      2 fe-0/1/0.0

```

show route forwarding-table vpls (Broadcast, unknown unicast, and multicast (BUM) hashing is enabled)

```

user@host> show route forwarding-table vpls
Routing table: green.vpls
VPLS:
Enabled protocols: BUM hashing
Destination      Type RtRef Next hop      Type Index  NhRef Netif
default          perm  0
lsi.1048832      intf  0
                  4.4.3.2      indr 1048574  4
                  Push 262145      621      2

ge-3/0/0.0
00:19:e2:25:d0:01/48 user  0                ucst  590      5 ge-2/3/9.0
0x30003/51       user  0                comp  627      2
ge-2/3/9.0       intf  0                ucst  590      5 ge-2/3/9.0
ge-3/1/3.0       intf  0                ucst  619      4 ge-3/1/3.0
0x30002/51       user  0                comp  600      2
0x30001/51       user  0                comp  597      2

```

show route forwarding-table vpls (Broadcast, unknown unicast, and multicast (BUM) hashing is enabled with MAC Statistics)

```

user@host> show route forwarding-table vpls
Routing table: green.vpls
VPLS:
Enabled protocols: BUM hashing, MAC Stats
Destination      Type RtRef Next hop      Type Index  NhRef Netif
default          perm  0         4.4.3.2      dscd   519      1
1si.1048834      intf  0         4.4.3.2      indr  1048574  4
                4.4.3.2      Push 262145 592      2
ge-3/0/0.0
00:19:e2:25:d0:01/48 user  0         ucst   590      5 ge-2/3/9.0
0x30003/51       user  0         comp   630      2
ge-2/3/9.0       intf  0         ucst   590      5 ge-2/3/9.0
ge-3/1/3.0       intf  0         ucst   591      4 ge-3/1/3.0
0x30002/51       user  0         comp   627      2
0x30001/51       user  0         comp   624      2

```

show route forwarding-table family vpls extensive

```

user@host> show route forwarding-table family vpls extensive
Routing table: green.vpls [Index 2]
VPLS:

Destination: default
Route type: dynamic
Route reference: 0
Flags: sent to PFE
Next-hop type: flood
Next-hop type: unicast
Next-hop interface: fe-0/1/3.0
Next-hop type: unicast
Next-hop interface: fe-0/1/2.0
Route interface-index: 72
Index: 289 Reference: 1
Index: 291 Reference: 3
Index: 290 Reference: 3

Destination: default
Route type: permanent
Route reference: 0
Flags: none
Next-hop type: discard
Route interface-index: 0
Index: 341 Reference: 1

Destination: fe-0/1/2.0
Route type: dynamic
Route reference: 0
Flags: sent to PFE
Next-hop type: flood
Next-hop type: indirect
Next-hop type: Push 800016
Next-hop interface: at-1/0/1.0
Next-hop type: indirect
Next hop: 10.31.3.2
Next-hop type: Push 800000
Next-hop interface: fe-0/1/1.0
Next-hop type: unicast
Next-hop interface: fe-0/1/3.0
Route interface-index: 69
Index: 293 Reference: 1
Index: 363 Reference: 4
Index: 301 Reference: 5
Index: 291 Reference: 3

Destination: fe-0/1/3.0
Route type: dynamic
Route reference: 0
Flags: sent to PFE
Next-hop type: flood
Route interface-index: 70
Index: 292 Reference: 1

```

```

Next-hop type: indirect          Index: 363      Reference: 4
Next-hop type: Push 800016
Next-hop interface: at-1/0/1.0
Next-hop type: indirect          Index: 301      Reference: 5
Next hop: 10.31.3.2
Next-hop type: Push 800000
Next-hop interface: fe-0/1/1.0
Next-hop type: unicast           Index: 290      Reference: 3
Next-hop interface: fe-0/1/2.0

Destination: 10:00:00:01:01:01/48
Route type: dynamic
Route reference: 0                Route interface-index: 70
Flags: sent to PFE, prefix load balance
Next-hop type: unicast           Index: 291      Reference: 3
Next-hop interface: fe-0/1/3.0
Route used as destination:
  Packet count:      6640    Byte count:      675786
Route used as source
  Packet count:      6894    Byte count:      696424

Destination: 10:00:00:01:01:04/48
Route type: dynamic
Route reference: 0                Route interface-index: 69
Flags: sent to PFE, prefix load balance
Next-hop type: unicast           Index: 290      Reference: 3
Next-hop interface: fe-0/1/2.0
Route used as destination:
  Packet count:      96     Byte count:      8079
Route used as source:
  Packet count:      296    Byte count:      24955

Destination: 10:00:00:01:03:05/48
Route type: dynamic
Route reference: 0                Route interface-index: 74
Flags: sent to PFE, prefix load balance
Next-hop type: indirect          Index: 301      Reference: 5
Next hop: 10.31.3.2
Next-hop type: Push 800000
Next-hop interface: fe-0/1/1.0

```

show route forwarding-table table default

```

user@host> show route forwarding-table table default
Routing table: default.inet
Internet:
Destination      Type RtRef Next hop          Type Index NhRef Netif
default          perm  0
0.0.0.0/32       perm  0
10.0.60.0/30     user  0 10.0.60.13          ucst  713  5 fe-0/1/3.0
10.0.60.12/30    intf  0
10.0.60.12/32    dest  0 10.0.60.12          recv  686  1 fe-0/1/3.0
10.0.60.13/32    dest  0 0:5:85:8b:bc:22     ucst  713  5 fe-0/1/3.0
10.0.60.14/32    intf  0 10.0.60.14          locl  687  2
10.0.60.14/32    dest  0 10.0.60.14          locl  687  2
10.0.60.15/32    dest  0 10.0.60.15          bcst  685  1 fe-0/1/3.0
10.0.67.12/30    user  0 10.0.60.13          ucst  713  5 fe-0/1/3.0
10.0.80.0/30     ifdn  0 ff.3.0.21          ucst  676  1 so-0/0/1.0
10.0.80.0/32     dest  0 10.0.80.0          recv  678  1 so-0/0/1.0
10.0.80.2/32     user  0
10.0.80.2/32     intf  0 10.0.80.2          locl  675  1

```

```

10.0.80.3/32      dest      0 10.0.80.3      bcst    677      1 so-0/0/1.0
10.0.90.12/30     intf      0                rslv    684      1 fe-0/1/0.0
10.0.90.12/32     dest      0 10.0.90.12     recv    682      1 fe-0/1/0.0
10.0.90.14/32     intf      0 10.0.90.14     locl    683      2
10.0.90.14/32     dest      0 10.0.90.14     locl    683      2
10.0.90.15/32     dest      0 10.0.90.15     bcst    681      1 fe-0/1/0.0
10.5.0.0/16       user      0 192.168.187.126 ucst    324     15 fxp0.0
10.10.0.0/16      user      0 192.168.187.126 ucst    324     15 fxp0.0
10.13.10.0/23     user      0 192.168.187.126 ucst    324     15 fxp0.0
10.84.0.0/16      user      0 192.168.187.126 ucst    324     15 fxp0.0
10.150.0.0/16     user      0 192.168.187.126 ucst    324     15 fxp0.0
10.157.64.0/19    user      0 192.168.187.126 ucst    324     15 fxp0.0
10.209.0.0/16     user      0 192.168.187.126 ucst    324     15 fxp0.0

```

...

Routing table: default.iso

ISO:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	60	1	

Routing table: default.inet6

Internet6:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	44	1	
::/128	perm	0		dscd	42	1	
ff00::/8	perm	0		mdsc	43	1	
ff02::1/128	perm	0	ff02::1	mcst	39	1	

Routing table: default.mpls

MPLS:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		dscd	50	1	

show route forwarding-table table logical-system-name/routing-instance-name

```
user@host> show route forwarding-table table R4/vpn-red
```

Logical system: R4

Routing table: vpn-red.inet

Internet:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	563	1	
0.0.0.0/32	perm	0		dscd	561	2	
1.0.0.1/32	user	0		dscd	561	2	
2.0.2.0/24	intf	0		rslv	771	1	ge-1/2/0.3
2.0.2.0/32	dest	0	2.0.2.0	recv	769	1	ge-1/2/0.3
2.0.2.1/32	intf	0	2.0.2.1	locl	770	2	
2.0.2.1/32	dest	0	2.0.2.1	locl	770	2	
2.0.2.2/32	dest	0	0.4.80.3.0.1b.c0.d5.e4.bd.0.1b.c0.d5.e4.bc.8.0	ucst	789	1	ge-1/2/0.3
2.0.2.255/32	dest	0	2.0.2.255	bcst	768	1	ge-1/2/0.3
224.0.0.0/4	perm	1		mdsc	562	1	
224.0.0.1/32	perm	0	224.0.0.1	mcst	558	1	
255.255.255.255/32	perm	0		bcst	559	1	

Logical system: R4

Routing table: vpn-red.iso

ISO:

Destination	Type	RtRef	Next hop	Type	Index	NhRef	Netif
default	perm	0		rjct	608	1	


```

Logical system: R4
Routing table: vpn-red.inet6
Internet6:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm  0              rjct  708    1
::/128           perm  0              dscd  706    1
ff00::/8         perm  0              mdsc  707    1
ff02::1/128     perm  0 ff02::1      mcst  704    1

```

```

Logical system: R4
Routing table: vpn-red.mpls
MPLS:
Destination      Type RtRef Next hop      Type Index NhRef Netif
default          perm  0              dscd  638

```

show route forwarding-table vpn

```

user@host> show route forwarding-table vpn VPN-A
Routing table:: VPN-A.inet
Internet:
Destination      Type RtRef Nexthop      Type Index NhRef Netif
default          perm  0              rjct   4    4
10.39.10.20/30   intf  0 ff.3.0.21      ucst   40    1
so-0/0/0.0
10.39.10.21/32   intf  0 10.39.10.21     locl   36    1
10.255.14.172/32 user  0              ucst   69    2
so-0/0/0.0
10.255.14.175/32 user  0              indr   81    3
Push 100004, Push
100004(top) so-1/0/0.0
224.0.0.0/4      perm  2              mdsc   5    3
224.0.0.1/32     perm  0 224.0.0.1      mcst   1    8
224.0.0.5/32     user  1 224.0.0.5      mcst   1    8
255.255.255.255/32 perm  0              bcst   2    3

```

show route forwarding-table interface-name

Syntax	show route forwarding-table interface-name <i>interface-name</i> <detail extensive> <all>
Release Information	Command introduced in Junos OS Release 9.6.
Description	Display the interfaces in the Routing Engine's forwarding table.
Options	<p>none—Display information for the specified interface.</p> <p>detail extensive—(Optional) Display the specified level of output.</p> <p>all—(Optional) Display all interfaces in the routing table.</p>
Required Privilege Level	view
List of Sample Output	show route forwarding-table interface-name fe-0/1/1 on page 865 show route forwarding-table interface-name all on page 865 show route forwarding-table interface-name all detail on page 866
Output Fields	Table 135 on page 864 lists the output fields for the show route forwarding-table interface-name command. Output fields are listed in the approximate order in which they appear.

Table 135: show route forwarding-table interface-name Output Fields

Field Name	Field Description	Level of Output
Name	Name of the interface (for example fe-0/1/1 , lo0 , ae0 , and so on).	All levels
MTU	Interface's maximum transmission unit (MTU).	All levels
Afam	Configured address family (for example inet , tnp , inet6 , and so on).	detail extensive
Network	Network information: <ul style="list-style-type: none"> • <Link>—Physical interface, not a logical interface. • <PtoP>—Point-to-point network. • ipaddress—Network address. 	All levels
Address	Address of the interface. The address can be a MAC address, IPv4 address, IPv6 address, and so on.	All levels
IPkts	Number of packets received on the interface.	All levels
Ierr	Number of packets received on the interface with errors.	All levels
Opkts	Number of packets transmitted or sent from the interface.	All levels

Table 135: show route forwarding-table interface-name Output Fields (*continued*)

Field Name	Field Description	Level of Output
Oerr	Number of packets transmitted or sent from the interface with errors.	All levels
Coll	Number of packets that experienced collisions on the interface.	All levels

Sample Output

show route forwarding-table interface-name fe-0/1/1

```

user@host> show route forwarding-table interface-name fe-0/1/1
Name      Mtu Network      Address      IpKts Ierr      Opkts Oerr  Coll
fe-0/1/1  1514 <Link>      00.05.85.88.cc.20      0    0        0    0    0

```

show route forwarding-table interface-name all

```

user@host> show route forwarding-table interface-name all
Name      Mtu Network      Address      IpKts Ierr      Opkts Oerr  Coll
fxp0      1514 <Link>      00.a0.a5.56.03.83      180965    0      39907    0    0

  unit 0
  fxp1   1514 <Link>      02.00.00.00.00.04      33010497    0      30110800    0    0

  unit 0
    1500 192.168.187.0/ 192.168.187.10
  fxp1   1514 <Link>      02.00.00.00.00.04      33010497    0      30110800    0    0

  unit 0
    1500 10.0.0.0/8      10.0.0.1
    10.0.0.0/8      10.0.0.4
    128.0.0.0/2      128.0.0.1
    128.0.0.0/2      128.0.0.4
    1500 fe80::/64      fe80::200:ff:fe0
    fec0::/64      fec0::a:0:0:4
    1500      4
  lsi     1496 <Link>
  dsc     max <Link>
  lo0     max <Link>
                                0    0        0    0    0
                                8980    0      8980    0    0

  unit 0
    max 127.0.0.1/8      127.0.0.1
    192.168.0.1/8      192.168.0.1
  unit 16384 max 127.0.0.1/8      127.0.0.1
  unit 16385 max
  gre     max <Link>
  ipip    max <Link>
  tap     max <Link>
  pime    max <Link>
  pimd    max <Link>
  mtun    max <Link>
  so-0/0/0 4474 <Link>
                                1679900    0      1068611    0    0

  unit 0
  0       4470 <PtoP>      10.0.60.2
                                0    0        0    0
  so-0/0/1 4474 <Link>
                                0    0        0    0    0

  unit 0
  0       4470 <PtoP>      10.0.80.2
                                0    0        0    0
  so-0/0/2 4474 <Link>
                                0    0        0    0    0
  so-0/0/3 4474 <Link>
                                0    0        0    0    0
  fe-0/1/0 1514 <Link>      00.05.85.88.cc.1f      523120    0      623044    0    0

  unit 0
    1500 10.0.90.12/30 10.0.90.14
                                0    0        0    0

```

```

0
fe-0/1/1 1514 <Link> 00.05.85.88.cc.20 0 0 0 0 0
fe-0/1/2 1514 <Link> 00.05.85.88.cc.21 0 0 0 0 0

```

```
...
```

show route forwarding-table interface-name all detail

```

user@host> show route forwarding-table interface-name all detail
Name      Mtu AFam  Network      Address      Ipkts Ierr  Opkts
Oerr Coll
fxp0      1514      <Link> 00.a0.a5.56.03.83 181005 0 39948
0 0
unit 0    1500 inet  192.168.187.0/ 192.168.187.10
fxp1      1514      <Link> 02.00.00.00.00.04 33012676 0 30112468
0 0
unit 0    1500 inet  10.0.0.0/8     10.0.0.1
                        10.0.0.0/8     10.0.0.4
                        128.0.0.0/2    128.0.0.1
                        128.0.0.0/2    128.0.0.4
                        1500 inet6  fe80::/64      fe80::200:ff:fe0
                        fec0::/64      fec0::a:0:0:4
                        1500 tnp
lsi        1496      <Link>
dsc         max      <Link> 0 0 0
0 0
lo0         max      <Link> 8980 0 8980
0 0
unit 0      max inet  127.0.0.1/8    127.0.0.1
                        192.168.0.1/8  192.168.0.1
unit 16384 max inet  127.0.0.1/8    127.0.0.1
unit 16385 max inet
gre         max      <Link>
ipip        max      <Link>
tap         max      <Link>
pime        max      <Link>
pimd        max      <Link>
mtun        max      <Link>
so-0/0/0    4474      <Link> 1679980 0 1068661
0 0
unit 0      4470 inet  <PtoP> 10.0.60.2 0 0 0
0 0
...

```

show route hidden

Syntax	show route hidden <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display only hidden route information. A hidden route is unusable, even if it is the best path.
Options	<p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>Understanding Hidden Routes</i>
List of Sample Output	show route hidden on page 867 show route hidden detail on page 868 show route hidden extensive on page 868 show route hidden terse on page 868
Output Fields	For information about output fields, see the output field table for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route hidden

```

user@host> show route hidden
inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
127.0.0.1/32      [Direct/0] 04:26:38
                  > via lo0.0

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.5.5.5/32      [BGP/170] 03:44:10, localpref 100, from 10.4.4.4
                  AS path: 100 I
                  Unusable
10.12.1.0/24     [BGP/170] 03:44:10, localpref 100, from 10.4.4.4
                  AS path: 100 I
                  Unusable

```

```

10.12.80.4/30      [BGP/170] 03:44:10, localpref 100, from 10.4.4.4
                  AS path: I
                  Unusable
...

```

show route hidden detail

```

user@host> show route hidden detail

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
127.0.0.1/32 (1 entry, 0 announced)
    Direct Preference: 0
        Next hop type: Interface
        Next-hop reference count: 1
        Next hop: via lo0.0, selected
        State: <Hidden Martian Int>
        Local AS:      1
        Age: 4:27:37
        Task: IF
        AS path: I

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete

10.5.5.5/32 (1 entry, 0 announced)
    BGP Preference: 170/-101
        Route Distinguisher: 10.4.4.4:4
        Next hop type: Unusable
        Next-hop reference count: 6
        State: <Secondary Hidden Int Ext>
        Local AS:      1 Peer AS:      1
        Age: 3:45:09
        Task: BGP_1.10.4.4.4+2493
        AS path: 100 I
        Communities: target:1:999
        VPN Label: 100064
        Localpref: 100
        Router ID: 10.4.4.4
        Primary Routing Table bgp.13vpn.0

...

```

show route hidden extensive

The output for the **show route hidden extensive** command is identical to that of the **show route hidden detail** command. For sample output, see [show route hidden detail on page 868](#).

show route hidden terse

```

user@host> show route hidden terse

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
127.0.0.1/32      D  0                >100.0

```

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)

Restart Complete

+ = Active Route, - = Last Active, * = Both

A Destination	P Prf	Metric 1	Metric 2	Next hop	AS path
10.5.5.5/32	B 170	100		Unusable	100 I
10.12.1.0/24	B 170	100		Unusable	100 I
10.12.80.4/30	B 170	100		Unusable	I

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

Restart Complete

bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)

Restart Complete

+ = Active Route, - = Last Active, * = Both

A Destination	P Prf	Metric 1	Metric 2	Next hop	AS path
10.4.4.4:4:10.5.5.5/32	B 170	100		Unusable	100 I
10.4.4.4:4:10.12.1.0/24	B 170	100		Unusable	100 I
10.4.4.4:4:10.12.80.4/30	B 170	100		Unusable	I

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

show route inactive-path

Syntax	<code>show route inactive-path</code> <code><brief detail extensive terse></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches)	<code>show route inactive-path</code> <code><brief detail extensive terse></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display routes for destinations that have no active route. An inactive route is a route that was not selected as the best path.
Options	none —Display all inactive routes. brief detail extensive terse —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief . logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route inactive-path on page 870 show route inactive-path detail on page 871 show route inactive-path extensive on page 872 show route inactive-path terse on page 872
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route inactive-path

```
user@host> show route inactive-path

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.12.100.12/30      [OSPF/10] 03:57:28, metric 1
> via so-0/3/0.0

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.0/8          [Direct/0] 04:39:56
> via fxp1.0
```



```

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.12.80.0/30      [BGP/170] 04:38:17, localpref 100
                  AS path: 100 I
                  > to 10.12.80.1 via ge-6/3/2.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

show route inactive-path detail

```

user@host> show route inactive-path detail

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete

10.12.100.12/30 (2 entries, 1 announced)
  OSPF   Preference: 10
         Next-hop reference count: 1
         Next hop: via so-0/3/0.0, selected
         State: <Int>
         Inactive reason: Route Preference
         Local AS:      1
         Age: 3:58:24   Metric: 1
         Area: 0.0.0.0
         Task: OSPF
         AS path: I

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

10.0.0.0/8 (2 entries, 0 announced)
  Direct Preference: 0
         Next hop type: Interface
         Next-hop reference count: 1
         Next hop: via fxp1.0, selected
         State: <NotBest Int>
         Inactive reason: No difference
         Age: 4:40:52
         Task: IF
         AS path: I

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete

10.12.80.0/30 (2 entries, 1 announced)
  BGP   Preference: 170/-101
        Next-hop reference count: 6
        Source: 10.12.80.1

```

```

Next hop: 10.12.80.1 via ge-6/3/2.0, selected
State: <Ext>
Inactive reason: Route Preference
Peer AS: 100
Age: 4:39:13
Task: BGP_100.10.12.80.1+179
AS path: 100 I
Localpref: 100
Router ID: 10.0.0.0

```

show route inactive-path extensive

The output for the **show route inactive-path extensive** command is identical to that of the **show route inactive-path detail** command. For sample output, see [show route inactive-path detail on page 871](#).

show route inactive-path terse

```
user@host> show route inactive-path terse
```

```
inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
```

```
Restart Complete
```

```
+ = Active Route, - = Last Active, * = Both
```

A	Destination	P	Prf	Metric 1	Metric 2	Next hop	AS path
	10.12.100.12/30	0	10	1		>so-0/3/0.0	

```
private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
```

```
+ = Active Route, - = Last Active, * = Both
```

A	Destination	P	Prf	Metric 1	Metric 2	Next hop	AS path
	10.0.0.0/8	D	0			>fxp1.0	

```
red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
```

```
Restart Complete
```

```
+ = Active Route, - = Last Active, * = Both
```

A	Destination	P	Prf	Metric 1	Metric 2	Next hop	AS path
	10.12.80.0/30	B	170	100		>10.12.80.1	100 I

```
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

```
Restart Complete
```

```
mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
```

```
Restart Complete
```

```
bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
```

```
Restart Complete
```

```
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
```

```
Restart Complete
```

```
private1__inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route inactive-prefix

Syntax	show route inactive-prefix <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show route inactive-prefix <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display inactive route destinations in each routing table.
Options	<p>none—Display all inactive route destination.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show route inactive-prefix on page 873 show route inactive-prefix detail on page 873 show route inactive-prefix extensive on page 874 show route inactive-prefix terse on page 874
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route inactive-prefix

```

user@host> show route inactive-prefix

inet.0: 14 destinations, 14 routes (13 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

127.0.0.1/32          [Direct/0] 00:04:54
> via lo0.0

```

show route inactive-prefix detail

```

user@host> show route inactive-prefix detail

inet.0: 14 destinations, 14 routes (13 active, 0 holddown, 1 hidden)
127.0.0.1/32 (1 entry, 0 announced)
    Direct Preference: 0
    Next hop type: Interface

```

```
Next-hop reference count: 1
Next hop: via lo0.0, selected
State: <Hidden Martian Int>
Age: 4:51
Task: IF
AS path: I00:04:54
> via lo0.0
```

`show route inactive-prefix extensive`

The output for the `show route inactive-prefix extensive` command is identical to that of the `show route inactive-path detail` command. For sample output, see [show route inactive-prefix detail on page 873](#).

`show route inactive-prefix terse`

```
user@host> show route inactive-prefix terse
```

```
inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
```

A	Destination	P	Prf	Metric 1	Metric 2	Next hop	AS path
	127.0.0.1/32	D	0			>lo0.0	

show route instance

Syntax	<pre>show route instance <brief detail summary> <instance-name> <logical-system (all logical-system-name)> <operational></pre>
Syntax (EX Series Switches and QFX Series)	<pre>show route instance <brief detail summary> <instance-name> <operational></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.3 for the QFX Series.</p>
Description	Display routing instance information.
Options	<p>none—(Same as brief) Display standard information about all routing instances.</p> <p>brief detail summary—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief. (These options are not available with the operational keyword.)</p> <p>instance-name—(Optional) Display information for all routing instances whose name begins with this string (for example, cust1, cust11, and cust111 are all displayed when you run the show route instance cust1 command).</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>operational—(Optional) Display operational routing instances.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Transporting IPv6 Traffic Across IPv4 Using Filter-Based Tunneling</i> • <i>Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart</i>
List of Sample Output	<p>show route instance on page 876</p> <p>show route instance detail (Graceful Restart Complete) on page 877</p> <p>show route instance detail (Graceful Restart Incomplete) on page 879</p> <p>show route instance detail (VPLS Routing Instance) on page 880</p> <p>show route instance operational on page 881</p> <p>show route instance summary on page 881</p>
Output Fields	<p>Table 136 on page 876 lists the output fields for the show route instance command. Output fields are listed in the approximate order in which they appear.</p>

Table 136: show route instance Output Fields

Field Name	Field Description	Level of Output
Instance or <i>instance-name</i>	Name of the routing instance.	All levels
Operational Routing Instances	(operational keyword only) Names of all operational routing instances.	—
Type	Type of routing instance: forwarding , l2vpn , no-forwarding , vpls , virtual-router , or vrf .	All levels
State	State of the routing instance: active or inactive .	brief detail none
Interfaces	Name of interfaces belonging to this routing instance.	brief detail none
Restart State	Status of graceful restart for this instance: Pending or Complete .	detail
Path selection timeout	Maximum amount of time, in seconds, remaining until graceful restart is declared complete. The default is 300 .	detail
Tables	Tables (and number of routes) associated with this routing instance.	brief detail none
Route-distinguisher	Unique route distinguisher associated with this routing instance.	detail
Vrf-import	VPN routing and forwarding instance import policy name.	detail
Vrf-export	VPN routing and forwarding instance export policy name.	detail
Vrf-import-target	VPN routing and forwarding instance import target community name.	detail
Vrf-export-target	VPN routing and forwarding instance export target community name.	detail
Fast-reroute-priority	Fast reroute priority setting for a VPLS routing instance: high , medium , or low . The default is low .	detail
Restart State	Restart state: <ul style="list-style-type: none"> • Pending;protocol-name—List of protocols that have not yet completed graceful restart for this routing table. • Complete—All protocols have restarted for this routing table. 	detail
Primary rib	Primary table for this routing instance.	brief none summary
Active/holddown/hidden	Number of active, hold-down, and hidden routes.	All levels

Sample Output

show route instance

```

user@host> show route instance
Instance          Type
                  Primary RIB
Active/holddown/hidden

```

```

master          forwarding
    inet.0      16/0/1
    iso.0       1/0/0
    mpls.0      0/0/0
    inet6.0     2/0/0
    l2circuit.0 0/0/0
__juniper_private1__ forwarding
    __juniper_private1__.inet.0 12/0/0
    __juniper_private1__.inet6.0 1/0/0

```

show route instance detail (Graceful Restart Complete)

```

user@host> show route instance detail
master:
  Router ID: 10.255.14.176
  Type: forwarding      State: Active
  Restart State: Complete Path selection timeout: 300
  Tables:
    inet.0              : 17 routes (15 active, 0 holddown, 1 hidden)
    Restart Complete
    inet.3              : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    iso.0               : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
    mpls.0              : 19 routes (19 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l3vpn.0         : 10 routes (10 active, 0 holddown, 0 hidden)
    Restart Complete
    inet6.0             : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l2vpn.0         : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
  BGP-INET:
    Router ID: 10.69.103.1
    Type: vrf           State: Active
    Restart State: Complete Path selection timeout: 300
    Interfaces:
      t3-0/0/0.103
    Route-distinguisher: 10.255.14.176:103
    Vrf-import: [ BGP-INET-import ]
    Vrf-export: [ BGP-INET-export ]
    Tables:
      BGP-INET.inet.0    : 4 routes (4 active, 0 holddown, 0 hidden)
      Restart Complete
  BGP-L:
    Router ID: 10.69.104.1
    Type: vrf           State: Active
    Restart State: Complete Path selection timeout: 300
    Interfaces:
      t3-0/0/0.104
    Route-distinguisher: 10.255.14.176:104
    Vrf-import: [ BGP-L-import ]
    Vrf-export: [ BGP-L-export ]
    Tables:
      BGP-L.inet.0       : 4 routes (4 active, 0 holddown, 0 hidden)
      Restart Complete
      BGP-L.mpls.0       : 3 routes (3 active, 0 holddown, 0 hidden)
      Restart Complete
  L2VPN:
    Router ID: 0.0.0.0
    Type: l2vpn         State: Active

```

```
Restart State: Complete Path selection timeout: 300
Interfaces:
  t3-0/0/0.512
Route-distinguisher: 10.255.14.176:512
Vrf-import: [ L2VPN-import ]
Vrf-export: [ L2VPN-export ]
Tables:
  L2VPN.l2vpn.0          : 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
LDP:
Router ID: 10.69.105.1
Type: vrf                State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
  t3-0/0/0.105
Route-distinguisher: 10.255.14.176:105
Vrf-import: [ LDP-import ]
Vrf-export: [ LDP-export ]
Tables:
  LDP.inet.0             : 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
OSPF:
Router ID: 10.69.101.1
Type: vrf                State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
  t3-0/0/0.101
Route-distinguisher: 10.255.14.176:101
Vrf-import: [ OSPF-import ]
Vrf-export: [ OSPF-export ]
Vrf-import-target: [ target:11111
Tables:
  OSPF.inet.0            : 8 routes (7 active, 0 holddown, 0 hidden)
Restart Complete
RIP:
Router ID: 10.69.102.1
Type: vrf                State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
  t3-0/0/0.102
Route-distinguisher: 10.255.14.176:102
Vrf-import: [ RIP-import ]
Vrf-export: [ RIP-export ]
Tables:
  RIP.inet.0             : 6 routes (6 active, 0 holddown, 0 hidden)
Restart Complete
STATIC:
Router ID: 10.69.100.1
Type: vrf                State: Active
Restart State: Complete Path selection timeout: 300
Interfaces:
  t3-0/0/0.100
Route-distinguisher: 10.255.14.176:100
Vrf-import: [ STATIC-import ]
Vrf-export: [ STATIC-export ]
Tables:
  STATIC.inet.0          : 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
```


show route instance detail (Graceful Restart Incomplete)

```

user@host> show route instance detail
master:
  Router ID: 10.255.14.176
  Type: forwarding          State: Active
  Restart State: Pending    Path selection timeout: 300
  Tables:
    inet.0                  : 17 routes (15 active, 1 holddown, 1 hidden)
    Restart Pending: OSPF LDP
    inet.3                  : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: OSPF LDP
    iso.0                   : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
    mpls.0                  : 23 routes (23 active, 0 holddown, 0 hidden)
    Restart Pending: LDP VPN
    bgp.l3vpn.0             : 10 routes (10 active, 0 holddown, 0 hidden)
    Restart Pending: BGP VPN
    inet6.0                 : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l2vpn.0             : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Pending: BGP VPN
BGP-INET:
  Router ID: 10.69.103.1
  Type: vrf                 State: Active
  Restart State: Pending    Path selection timeout: 300
  Interfaces:
    t3-0/0/0.103
  Route-distinguisher: 10.255.14.176:103
  Vrf-import: [ BGP-INET-import ]
  Vrf-export: [ BGP-INET-export ]
  Tables:
    BGP-INET.inet.0        : 6 routes (5 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
BGP-L:
  Router ID: 10.69.104.1
  Type: vrf                 State: Active
  Restart State: Pending    Path selection timeout: 300
  Interfaces:
    t3-0/0/0.104
  Route-distinguisher: 10.255.14.176:104
  Vrf-import: [ BGP-L-import ]
  Vrf-export: [ BGP-L-export ]
  Tables:
    BGP-L.inet.0           : 6 routes (5 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
    BGP-L.mpls.0           : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
L2VPN:
  Router ID: 0.0.0.0
  Type: l2vpn               State: Active
  Restart State: Pending    Path selection timeout: 300
  Interfaces:
    t3-0/0/0.512
  Route-distinguisher: 10.255.14.176:512
  Vrf-import: [ L2VPN-import ]
  Vrf-export: [ L2VPN-export ]
  Tables:
    L2VPN.l2vpn.0          : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: VPN L2VPN
LDP:

```

```
Router ID: 10.69.105.1
Type: vrf                      State: Active
Restart State: Pending Path selection timeout: 300
Interfaces:
  t3-0/0/0.105
Route-distinguisher: 10.255.14.176:105
Vrf-import: [ LDP-import ]
Vrf-export: [ LDP-export ]
Tables:
  LDP.inet.0                   : 5 routes (4 active, 1 holddown, 0 hidden)
Restart Pending: OSPF LDP VPN

OSPF:
  Router ID: 10.69.101.1
  Type: vrf                      State: Active
  Restart State: Pending Path selection timeout: 300
  Interfaces:
    t3-0/0/0.101
  Route-distinguisher: 10.255.14.176:101
  Vrf-import: [ OSPF-import ]
  Vrf-export: [ OSPF-export ]
  Tables:
    OSPF.inet.0                : 8 routes (7 active, 1 holddown, 0 hidden)
Restart Pending: OSPF VPN

RIP:
  Router ID: 10.69.102.1
  Type: vrf                      State: Active
  Restart State: Pending Path selection timeout: 300
  Interfaces:
    t3-0/0/0.102
  Route-distinguisher: 10.255.14.176:102
  Vrf-import: [ RIP-import ]
  Vrf-export: [ RIP-export ]
  Tables:
    RIP.inet.0                  : 8 routes (6 active, 2 holddown, 0 hidden)
Restart Pending: RIP VPN

STATIC:
  Router ID: 10.69.100.1
  Type: vrf                      State: Active
  Restart State: Pending Path selection timeout: 300
  Interfaces:
    t3-0/0/0.100
  Route-distinguisher: 10.255.14.176:100
  Vrf-import: [ STATIC-import ]
  Vrf-export: [ STATIC-export ]
  Tables:
    STATIC.inet.0               : 4 routes (4 active, 0 holddown, 0 hidden)
Restart Pending: VPN
```

show route instance detail (VPLS Routing Instance)

```
user@host> show route instance detail test-vpls
test-vpls:
  Router ID: 0.0.0.0
  Type: vpls                      State: Active
  Interfaces:
    lsi.1048833
    lsi.1048832
    fe-0/1/0.513
  Route-distinguisher: 10.255.37.65:1
  Vrf-import: [ __vrf-import-test-vpls-internal__ ]
  Vrf-export: [ __vrf-export-test-vpls-internal__ ]
```

```

Vrf-import-target: [ target:300:1 ]
Vrf-export-target: [ target:300:1 ]
Fast-reroute-priority: high
Tables:
  test-vpls.l2vpn.0          : 3 routes (3 active, 0 holddown, 0 hidden)

```

show route instance operational

```

user@host> show route instance operational
Operational Routing Instances:

master
default

```

show route instance summary

```

user@host> show route instance summary

```

Instance	Type	Primary rib	Active/holddown/hidden
master	forwarding	inet.0	15/0/1
		iso.0	1/0/0
		mpls.0	35/0/0
		l3vpn.0	0/0/0
		inet6.0	2/0/0
		l2vpn.0	0/0/0
		l2circuit.0	0/0/0
BGP-INET	vrf	BGP-INET.inet.0	5/0/0
		BGP-INET.iso.0	0/0/0
		BGP-INET.inet6.0	0/0/0
BGP-L	vrf	BGP-L.inet.0	5/0/0
		BGP-L.iso.0	0/0/0
		BGP-L.mpls.0	4/0/0
		BGP-L.inet6.0	0/0/0
L2VPN	l2vpn	L2VPN.inet.0	0/0/0
		L2VPN.iso.0	0/0/0
		L2VPN.inet6.0	0/0/0
		L2VPN.l2vpn.0	2/0/0
LDP	vrf	LDP.inet.0	4/0/0
		LDP.iso.0	0/0/0
		LDP.mpls.0	0/0/0
		LDP.inet6.0	0/0/0
		LDP.l2circuit.0	0/0/0
OSPF	vrf	OSPF.inet.0	7/0/0
		OSPF.iso.0	0/0/0
		OSPF.inet6.0	0/0/0
RIP	vrf	RIP.inet.0	6/0/0
		RIP.iso.0	0/0/0
		RIP.inet6.0	0/0/0
STATIC	vrf	STATIC.inet.0	4/0/0
		STATIC.iso.0	0/0/0
		STATIC.inet6.0	0/0/0

show route label

Syntax	<code>show route label <i>label</i></code> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	<code>show route label <i>label</i></code> <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
Description	Display the routes based on a specified Multiprotocol Label Switching (MPLS) label value.
Options	<i>label</i> —Value of the MPLS label. brief detail extensive terse —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring Multipoint LDP In-Band Signaling for Point-to-Multipoint LSPs</i>
List of Sample Output	show route label terse on page 882 show route label on page 883 show route label detail on page 883 show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 883 show route label detail (Multipoint LDP with Multicast-Only Fast Reroute) on page 883 show route label extensive on page 884
Output Fields	For information about output fields, see the output field table for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route label terse

```
user@host> show route label 100016 terse

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1   Metric 2   Next hop      AS path
* 100016           V 170                >10.12.80.1
```

show route label

```

user@host> show route label 100016

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
100016          *[VPN/170] 03:25:41
                > to 10.12.80.1 via ge-6/3/2.0, Pop

```

show route label detail

```

user@host> show route label 100016 detail

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
100016 (1 entry, 1 announced)
    *VPN      Preference: 170
              Next-hop reference count: 2
              Source: 10.12.80.1
              Next hop: 10.12.80.1 via ge-6/3/2.0, selected
              Label operation: Pop
              State: <Active Int Ext>
              Local AS: 1
              Age: 3:23:31
              Task: BGP.0.0.0.0+179
              Announcement bits (1): 0-KRT
              AS path: 100 I
              Ref Cnt: 2

```

show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```

user@host> show route label 299872 detail

mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
    *LDP      Preference: 9
              Next hop type: Flood
              Next-hop reference count: 3
              Address: 0x9097d90
              Next hop: via vt-0/1/0.1
              Next-hop index: 661
              Label operation: Pop
              Address: 0x9172130
              Next hop: via so-0/0/3.0
              Next-hop index: 654
              Label operation: Swap 299872
              State: **Active Int>
              Local AS: 1001
              Age: 8:20      Metric: 1
              Task: LDP
              Announcement bits (1): 0-KRT
              AS path: I
              FECs bound to route: P2MP root-addr 10.255.72.166, grp 232.1.1.1,
src 192.168.142.2

```

show route label detail (Multipoint LDP with Multicast-Only Fast Reroute)

```

user@host> show route label 301568 detail

mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)

```

```

*LDP      Preference: 9
          Next hop type: Flood
          Address: 0x2735208
          Next-hop reference count: 3
          Next hop type: Router, Next hop index: 1397
          Address: 0x2735d2c
          Next-hop reference count: 3
          Next hop: 1.3.8.2 via ge-1/2/22.0
          Label operation: Pop
          Load balance label: None;
          Next hop type: Router, Next hop index: 1395
          Address: 0x2736290
          Next-hop reference count: 3
          Next hop: 1.3.4.2 via ge-1/2/18.0
          Label operation: Pop
          Load balance label: None;
          State: <Active Int AckRequest MulticastRPF>
          Local AS: 10
          Age: 54:05      Metric: 1
          Validation State: unverified
          Task: LDP
          Announcement bits (1): 0-KRT
          AS path: I
          FECs bound to route: P2MP root-addr 1.1.1.1, grp: 232.1.1.1, src:
192.168.219.11
          Primary Upstream : 1.1.1.3:0--1.1.1.2:0
            RPF Nexthops :
              ge-1/2/15.0, 1.2.94.1, Label: 301568, weight: 0x1
              ge-1/2/14.0, 1.2.3.1, Label: 301568, weight: 0x1
          Backup Upstream : 1.1.1.3:0--1.1.1.6:0
            RPF Nexthops :
              ge-1/2/20.0, 1.2.96.1, Label: 301584, weight: 0xffff
              ge-1/2/19.0, 1.3.6.1, Label: 301584, weight: 0xffff

```

show route label extensive

The output for the `show route label extensive` command is identical to that of the `show route label detail` command. For sample output, see [show route label detail on page 883](#).

show route label-switched-path

Syntax	<code>show route label-switched-path <i>path-name</i></code> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	<code>show route label-switched-path <i>path-name</i></code> <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.5 for EX Series switches.
Description	Display the routes used in an MPLS label-switched path (LSP).
Options	brief detail extensive terse —(Optional) Display the specified level of output. <i>path-name</i> —LSP tunnel name. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route label-switched-path on page 885
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route label-switched-path

```

user@host> show route label-switched-path sf-to-ny
inet.0: 29 destinations, 29 routes (29 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.1/32          [MPLS/7] 00:00:06, metric 0
> to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
3.3.3.3/32          * [MPLS/7] 00:00:06, metric 0
> to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2.2.2.2/32          * [MPLS/7] 00:00:06, metric 0
> to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
4.4.4.4/32          * [MPLS/7] 00:00:06, metric 0
> to 111.222.1.9 via s0-0/0/0, label-switched-path abc
> to 111.222.1.9 via s0-0/0/0, label-switched-path xyz
> to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
111.222.1.9/32      [MPLS/7] 00:00:06, metric 0
> to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny

```

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

mpls.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

show route localization

Syntax	show route localization
Release Information	Command introduced in Junos OS Release 11.4 for T-Series routers. Command introduced in Junos OS Release 12.3 for MX Series routers.
Description	(T320, T640, and T1600 routers only) Display route localization details.
Options	detail —Display detailed output.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> <i>Example: Configuring Packet Forwarding Engine FIB Localization</i>
Output Fields	Table 137 on page 887 lists the output fields for the show route localization command. Output fields are listed in the approximate order in which they appear.

Table 137: show route localization Output Fields

Field Name	Field Description
FIB-local	FPCs configured as FIB-local.
FIB-remote	FPCs configured as FIB-remote.
Normal	FPCs neither configured as FIB-local or FIB-remote .
Protocols	IPv4 (inet) or IPv6 (inet6) traffic configured for route localization.

Sample Output

```

user@R0> show route localization
FIB localization ready FPCs (and FIB-local Forwarding Engine addresses)
  FIB-local:  FPC2(4,5)
  FIB-remote: FPC0, FPC1
  Normal:     FPC3, FPC4, FPC5, FPC6, FPC7

user@R0> show route localization detail
FIB localization ready FPCs (and FIB-local Forwarding Engine addresses)
  FIB-local:  FPC2(4,5)
  FIB-remote: FPC0, FPC1
  Normal:     FPC3, FPC4, FPC5, FPC6, FPC7
FIB localization configuration
  Protocols:  inet, inet6
  FIB-local:  FPC2
  FIB-remote: FPC0, FPC1
Forwarding Engine addresses
  FPC0: 1
  FPC1: 2
  FPC2: 4, 5

```

FPC3: 6
FPC4: 8
FPC5: 11
FPC6: 13
FPC7: 15

show route martians

Syntax	show route martians <logical-system (all <i>logical-system-name</i>)> <table <i>routing-table-name</i> >
Syntax (EX Series Switches)	show route martians <table <i>routing-table-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the martian (invalid and ignored) entries associated with each routing table.
Options	<p>none—Display standard information about route martians for all routing tables.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>table <i>routing-table-name</i>—(Optional) Display information about route martians for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the show route martians table inet command).</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Martian Addresses</i>
List of Sample Output	show route martians on page 889
Output Fields	Table 138 on page 889 lists the output fields for the show route martians command. Output fields are listed in the approximate order in which they appear

Table 138: show route martians Output Fields

Field Name	Field Description
<i>table-name</i>	Name of the route table in which the route martians reside.
<i>destination-prefix</i>	Route destination.
<i>match value</i>	Route match parameter.
<i>status</i>	Status of the route: allowed or disallowed .

Sample Output

show route martians

```
user@host> show route martians
```

```
inet.0:
    0.0.0.0/0 exact -- allowed
    0.0.0.0/8 orlonger -- disallowed
    127.0.0.0/8 orlonger -- disallowed
    192.0.0.0/24 orlonger -- disallowed
    240.0.0.0/4 orlonger -- disallowed
    224.0.0.0/4 exact -- disallowed
    224.0.0.0/24 exact -- disallowed

inet.1:
    0.0.0.0/0 exact -- allowed
    0.0.0.0/8 orlonger -- disallowed
    127.0.0.0/8 orlonger -- disallowed
    192.0.0.0/24 orlonger -- disallowed
    240.0.0.0/4 orlonger -- disallowed

inet.2:
    0.0.0.0/0 exact -- allowed
    0.0.0.0/8 orlonger -- disallowed
    127.0.0.0/8 orlonger -- disallowed
    192.0.0.0/24 orlonger -- disallowed
    240.0.0.0/4 orlonger -- disallowed
    224.0.0.0/4 exact -- disallowed
    224.0.0.0/24 exact -- disallowed

inet.3:
    0.0.0.0/0 exact -- allowed
    0.0.0.0/8 orlonger -- disallowed
    127.0.0.0/8 orlonger -- disallowed
    192.0.0.0/24 orlonger -- disallowed
    240.0.0.0/4 orlonger -- disallowed
    224.0.0.0/4 exact -- disallowed
    224.0.0.0/24 exact -- disallowed

...

inet6.0:
    ::1/128 exact -- disallowed
    ff00::/8 exact -- disallowed
    ff02::/16 exact -- disallowed

inet6.1:
    ::1/128 exact -- disallowed

inet6.2:
    ::1/128 exact -- disallowed
    ff00::/8 exact -- disallowed
    ff02::/16 exact -- disallowed

inet6.3:
    ::1/128 exact -- disallowed
    ff00::/8 exact -- disallowed
    ff02::/16 exact -- disallowed

...
```

show route match-prefix

Syntax	<code>show route match-prefix <i>match-prefix</i>;</code>
Release Information	Command introduced in Junos OS Release 11.4.
Description	Allows you to search for routes using regular expressions based on the extended (modern) regular expressions as defined in POSIX 1003.2.
Options	<i>match-prefix</i> —Regular expression to match formatted prefix.
Additional Information	
Required Privilege Level	view
Related Documentation	<i>Regular Expressions for Allowing and Denying Junos OS Operational Mode Commands</i>
List of Sample Output	show route match-prefix *:10.255.2.200:6:* (Show all routes matching route distributor 10.255.2.200:6) on page 891 show route match-prefix 7* (Show all mvpn type-7 routes) on page 891 show route match-prefix *:224.* (Show all routes matching group 224/4) on page 891
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

[show route match-prefix *:10.255.2.200:6:* \(Show all routes matching route distributor 10.255.2.200:6\)](#)

```
user@host> show route match-prefix *:10.255.2.200:6:*
```

[show route match-prefix 7* \(Show all mvpn type-7 routes\)](#)

```
user@host> show route table blue.mvpn.0 match-prefix 7*
Paste
router command output here
```

[show route match-prefix *:224.* \(Show all routes matching group 224/4\)](#)

```
user@host> show route match-prefix *:224.*
```

show route next-hop

Syntax	<code>show route next-hop <i>next-hop</i></code> <code><brief detail extensive terse></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches)	<code>show route next-hop <i>next-hop</i></code> <code><brief detail extensive terse></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the entries in the routing table that are being sent to the specified next-hop address.
Options	brief detail extensive terse —(Optional) Display the specified level of output. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. <i>next-hop</i> —Next-hop address.
Required Privilege Level	view
List of Sample Output	show route next-hop on page 892 show route next-hop detail on page 893 show route next-hop extensive on page 895 show route next-hop terse on page 896
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route next-hop

```
user@host> show route next-hop 192.168.71.254

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16      *[Static/5] 06:26:25
                  > to 192.168.71.254 via fxp0.0
10.209.0.0/16    *[Static/5] 06:26:25
                  > to 192.168.71.254 via fxp0.0
172.16.0.0/12    *[Static/5] 06:26:25
                  > to 192.168.71.254 via fxp0.0
192.168.0.0/16   *[Static/5] 06:26:25
                  > to 192.168.71.254 via fxp0.0
192.168.102.0/23 *[Static/5] 06:26:25
                  > to 192.168.71.254 via fxp0.0
```

```

207.17.136.0/24    *[Static/5] 06:26:25
                  > to 192.168.71.254 via fxp0.0
207.17.136.192/32 *[Static/5] 06:26:25
                  > to 192.168.71.254 via fxp0.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

show route next-hop detail

```

user@host> show route next-hop 192.168.71.254 detail

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
Restart Complete
10.10.0.0/16 (1 entry, 1 announced)
    *Static Preference: 5
        Next-hop reference count: 36
        Next hop: 192.168.71.254 via fxp0.0, selected
        State: <Active NoReadvrt Int Ext>
        Local AS:    1
        Age: 6:27:41
        Task: RT
        Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
        AS path: I

10.209.0.0/16 (1 entry, 1 announced)
    *Static Preference: 5
        Next-hop reference count: 36
        Next hop: 192.168.71.254 via fxp0.0, selected
        State: <Active NoReadvrt Int Ext>
        Local AS:    1
        Age: 6:27:41
        Task: RT
        Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
        AS path: I

172.16.0.0/12 (1 entry, 1 announced)
    *Static Preference: 5
        Next-hop reference count: 36
        Next hop: 192.168.71.254 via fxp0.0, selected
        State: <Active NoReadvrt Int Ext>
        Local AS:    1
        Age: 6:27:41
        Task: RT
        Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
        AS path: I

192.168.0.0/16 (1 entry, 1 announced)

```

```
*Static Preference: 5
  Next-hop reference count: 36
  Next hop: 192.168.71.254 via fxp0.0, selected
  State: <Active NoReadvrt Int Ext>
  Local AS: 1
  Age: 6:27:41
  Task: RT
  Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
  AS path: I

192.168.102.0/23 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

207.17.136.0/24 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

207.17.136.192/32 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 36
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 1
    Age: 6:27:41
    Task: RT
    Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
    AS path: I

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```


show route next-hop extensive

```
user@host> show route next-hop 192.168.71.254 extensive
```

```
inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
```

```
10.10.0.0/16 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
```

```
*Static Preference: 5
```

```
Next-hop reference count: 22
```

```
Next hop: 192.168.71.254 via fxp0.0, selected
```

```
State: <Active NoReadvrt Int Ext>
```

```
Local AS: 69
```

```
Age: 2:02:28
```

```
Task: RT
```

```
Announcement bits (1): 0-KRT
```

```
AS path: I
```

```
10.209.0.0/16 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
```

```
*Static Preference: 5
```

```
Next-hop reference count: 22
```

```
Next hop: 192.168.71.254 via fxp0.0, selected
```

```
State: <Active NoReadvrt Int Ext>
```

```
Local AS: 69
```

```
Age: 2:02:28
```

```
Task: RT
```

```
Announcement bits (1): 0-KRT
```

```
AS path: I
```

```
172.16.0.0/12 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 172.16.0.0/12 -> {192.168.71.254}
```

```
*Static Preference: 5
```

```
Next-hop reference count: 22
```

```
Next hop: 192.168.71.254 via fxp0.0, selected
```

```
State: <Active NoReadvrt Int Ext>
```

```
Local AS: 69
```

```
Age: 2:02:28
```

```
Task: RT
```

```
Announcement bits (1): 0-KRT
```

```
AS path: I
```

```
192.168.0.0/16 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 192.168.0.0/16 -> {192.168.71.254}
```

```
*Static Preference: 5
```

```
Next-hop reference count: 22
```

```
Next hop: 192.168.71.254 via fxp0.0, selected
```

```
State: <Active NoReadvrt Int Ext>
```

```
Local AS: 69
```

```
Age: 2:02:28
```

```
Task: RT
```

```
Announcement bits (1): 0-KRT
```

```
AS path: I
```

```
192.168.102.0/23 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 192.168.102.0/23 -> {192.168.71.254}
```

```
*Static Preference: 5
```

```

Next-hop reference count: 22
Next hop: 192.168.71.254 via fxp0.0, selected
State: <Active NoReadvrt Int Ext>
Local AS: 69
Age: 2:02:28
Task: RT
Announcement bits (1): 0-KRT
AS path: I

207.17.136.0/24 (1 entry, 1 announced)
TSI:
KRT in-kerne1 207.17.136.0/24 -> {192.168.71.254}
*Static Preference: 5
Next-hop reference count: 22
Next hop: 192.168.71.254 via fxp0.0, selected
State: <Active NoReadvrt Int Ext>
Local AS: 69
Age: 2:02:28
Task: RT
Announcement bits (1): 0-KRT
AS path: I

207.17.136.192/32 (1 entry, 1 announced)
TSI:
KRT in-kerne1 207.17.136.192/32 -> {192.168.71.254}
*Static Preference: 5
Next-hop reference count: 22
Next hop: 192.168.71.254 via fxp0.0, selected
State: <Active NoReadvrt Int Ext>
Local AS: 69
Age: 2:02:28
Task: RT
Announcement bits (1): 0-KRT
AS path: I

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

green.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

red.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

```

show route next-hop terse

```

user@host> show route next-hop 192.168.71.254 terse

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1   Metric 2   Next hop      AS path
* 10.10.0.0/16     S  5          0          0   >192.168.71.254
* 10.209.0.0/16    S  5          0          0   >192.168.71.254
* 172.16.0.0/12    S  5          0          0   >192.168.71.254

```

```
* 192.168.0.0/16      S   5                >192.168.71.254
* 192.168.102.0/23   S   5                >192.168.71.254
* 207.17.136.0/24    S   5                >192.168.71.254
* 207.17.136.192/32  S   5                >192.168.71.254

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route no-community

Syntax	show route no-community <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show route no-community <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the route entries in each routing table that are not associated with any community.
Options	none —(Same as brief) Display the route entries in each routing table that are not associated with any community. brief detail extensive terse —(Optional) Display the specified level of output. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route no-community on page 898 show route no-community detail on page 899 show route no-community extensive on page 899 show route no-community terse on page 900
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route no-community

```

user@host> show route no-community
inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16      *[Static/5] 00:36:27
                  > to 192.168.71.254 via fxp0.0
10.209.0.0/16    *[Static/5] 00:36:27
                  > to 192.168.71.254 via fxp0.0
10.255.71.52/32  *[Direct/0] 00:36:27
                  > via lo0.0
10.255.71.63/32  *[OSPF/10] 00:04:39, metric 1
                  > to 35.1.1.2 via ge-3/1/0.0
10.255.71.64/32  *[OSPF/10] 00:00:08, metric 2
                  > to 35.1.1.2 via ge-3/1/0.0
10.255.71.240/32 *[OSPF/10] 00:05:04, metric 2
                  via so-0/1/2.0

```

```

> via so-0/3/2.0
10.255.71.241/32  * [OSPF/10] 00:05:14, metric 1
> via so-0/1/2.0
10.255.71.242/32  * [OSPF/10] 00:05:19, metric 1
> via so-0/3/2.0
12.1.1.0/24      * [OSPF/10] 00:05:14, metric 2
> via so-0/3/2.0
14.1.1.0/24      * [OSPF/10] 00:00:08, metric 3
> to 35.1.1.2 via ge-3/1/0.0
                  via so-0/1/2.0
                  via so-0/3/2.0
16.1.1.0/24      * [OSPF/10] 00:05:14, metric 2
> via so-0/1/2.0
.....

```

show route no-community detail

```

user@host> show route no-community detail

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 38:08
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.209.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 38:08
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

....

```

show route no-community extensive

```

user@host> show route no-community extensive

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 2:03:33
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.209.0.0/16 (1 entry, 1 announced)
TSI:

```

```

KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Local AS: 69
    Age: 2:03:33
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

```

show route no-community terse

```
user@host> show route no-community terse
```

```

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

```

A	Destination	P	Prf	Metric 1	Metric 2	Next hop	AS path
*	10.10.0.0/16	S	5			>192.168.71.254	
*	10.209.0.0/16	S	5			>192.168.71.254	
*	10.255.71.52/32	D	0			>lo0.0	
*	10.255.71.63/32	0	10	1		>35.1.1.2	
*	10.255.71.64/32	0	10	2		>35.1.1.2	
*	10.255.71.240/32	0	10	2		so-0/1/2.0	
						>so-0/3/2.0	
*	10.255.71.241/32	0	10	1		>so-0/1/2.0	
*	10.255.71.242/32	0	10	1		>so-0/3/2.0	
*	12.1.1.0/24	0	10	2		>so-0/3/2.0	
*	14.1.1.0/24	0	10	3		>35.1.1.2	
						so-0/1/2.0	
						so-0/3/2.0	
*	16.1.1.0/24	0	10	2		>so-0/1/2.0	
...							

show route output

Syntax	show route output (address <i>ip-address</i> interface <i>interface-name</i>) <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show route output (address <i>ip-address</i> interface <i>interface-name</i>) <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Display the entries in the routing table learned through static routes and interior gateway protocols that are to be sent out the interface with either the specified IP address or specified name.</p> <p>To view routes advertised to a neighbor or received from a neighbor for the BGP protocol, use the show route advertising-protocol bgp and show route receive-protocol bgp commands instead.</p>
Options	<p>address <i>ip-address</i>—Display entries in the routing table that are to be sent out the interface with the specified IP address.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>interface <i>interface-name</i>—Display entries in the routing table that are to be sent out the interface with the specified name.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show route output address on page 902 show route output address detail on page 902 show route output address extensive on page 903 show route output address terse on page 903 show route output interface on page 903 show route output interface detail on page 904 show route output interface extensive on page 904 show route output interface terse on page 904
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route output address

```
user@host> show route output address 36.1.1.1/24

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

36.1.1.0/24          *[Direct/0] 00:19:56
                    > via so-0/1/2.0
                    [OSPF/10] 00:19:55, metric 1
                    > via so-0/1/2.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route output address detail

```
user@host> show route output address 36.1.1.1 detail

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
36.1.1.0/24 (2 entries, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via so-0/1/2.0, selected
    State: <Active Int>
    Age: 23:00
    Task: IF
    AS path: I
  OSPF Preference: 10
    Next-hop reference count: 1
    Next hop: via so-0/1/2.0, selected
    State: <Int>
    Inactive reason: Route Preference
    Age: 22:59      Metric: 1
    Area: 0.0.0.0
    Task: OSPF
    AS path: I

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```


show route output address extensive

The output for the **show route output address extensive** command is identical to that of the **show route output address detail** command. For sample output, see [show route output address detail on page 902](#).

show route output address terse

```
user@host> show route output address 36.1.1.1 terse

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1   Metric 2   Next hop      AS path
* 36.1.1.0/24      D   0                >so-0/1/2.0
                   O  10              1         >so-0/1/2.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route output interface

```
user@host> show route output interface so-0/1/2.0

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.71.240/32   *[OSPF/10] 00:13:00, metric 2
                   via so-0/1/2.0
                   > via so-0/3/2.0
10.255.71.241/32   *[OSPF/10] 00:13:10, metric 1
                   > via so-0/1/2.0
14.1.1.0/24        *[OSPF/10] 00:05:11, metric 3
                   to 35.1.1.2 via ge-3/1/0.0
                   > via so-0/1/2.0
                   via so-0/3/2.0
16.1.1.0/24        *[OSPF/10] 00:13:10, metric 2
                   > via so-0/1/2.0
36.1.1.0/24        *[Direct/0] 00:13:21
                   > via so-0/1/2.0
                   [OSPF/10] 00:13:20, metric 1
                   > via so-0/1/2.0

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route output interface detail

```
user@host> show route output interface so-0/1/2.0 detail
```

```
inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
10.255.71.240/32 (1 entry, 1 announced)
    *OSPF    Preference: 10
              Next-hop reference count: 2
              Next hop: via so-0/1/2.0
              Next hop: via so-0/3/2.0, selected
              State: <Active Int>
              Age: 14:52      Metric: 2
              Area: 0.0.0.0
              Task: OSPF
              Announcement bits (1): 0-KRT
              AS path: I

10.255.71.241/32 (1 entry, 1 announced)
    *OSPF    Preference: 10
              Next-hop reference count: 4
              Next hop: via so-0/1/2.0, selected
              State: <Active Int>
              Age: 15:02      Metric: 1
              Area: 0.0.0.0
              Task: OSPF
              Announcement bits (1): 0-KRT
              AS path: I

...
```

show route output interface extensive

The output for the **show route output interface extensive** command is identical to that of the **show route output interface detail** command. For sample output, see [show route output interface detail on page 904](#).

show route output interface terse

```
user@host> show route output interface so-0/1/2.0 terse
```

```
inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* 10.255.71.240/32  0 10      2              so-0/1/2.0
                        >so-0/3/2.0
* 10.255.71.241/32  0 10      1              >so-0/1/2.0
* 14.1.1.0/24       0 10      3              35.1.1.2
                        >so-0/1/2.0
                        so-0/3/2.0
* 16.1.1.0/24       0 10      2              >so-0/1/2.0
* 36.1.1.0/24       D  0
                        0 10      1              >so-0/1/2.0
                        >so-0/1/2.0

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
```

```
private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route protocol

Syntax	<code>show route protocol <i>protocol</i></code> <code><brief detail extensive terse></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches)	<code>show route protocol <i>protocol</i></code> <code><brief detail extensive terse></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. ospf2 and ospf3 options introduced in Junos OS Release 9.2. ospf2 and ospf3 options introduced in Junos OS Release 9.2 for EX Series switches. flow option introduced in Junos OS Release 10.0. flow option introduced in Junos OS Release 10.0 for EX Series switches.
Description	Display the route entries in the routing table that were learned from a particular protocol.
Options	brief detail extensive terse —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief . logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. <i>protocol</i> —Protocol from which the route was learned: <ul style="list-style-type: none">• access—Access route for use by DHCP application• access-internal—Access-internal route for use by DHCP application• aggregate—Locally generated aggregate route• arp—Route learned through the Address Resolution Protocol• atmvpn—Asynchronous Transfer Mode virtual private network• bgp—Border Gateway Protocol• ccc—Circuit cross-connect• direct—Directly connected route• dvmrp—Distance Vector Multicast Routing Protocol• esis—End System-to-Intermediate System• flow—Locally defined flow-specification route• frr—Precomputed protection route or backup route used when a link goes down• isis—Intermediate System-to-Intermediate System• ldp—Label Distribution Protocol• l2circuit—Layer 2 circuit• l2vpn—Layer 2 virtual private network

- **local**—Local address
- **mpls**—Multiprotocol Label Switching
- **msdp**—Multicast Source Discovery Protocol
- **ospf**—Open Shortest Path First versions 2 and 3
- **ospf2**—Open Shortest Path First versions 2 only
- **ospf3**—Open Shortest Path First version 3 only
- **pim**—Protocol Independent Multicast
- **rip**—Routing Information Protocol
- **ripng**—Routing Information Protocol next generation
- **rsvp**—Resource Reservation Protocol
- **rtarget**—Local route target virtual private network
- **static**—Statically defined route
- **tunnel**—Dynamic tunnel
- **vpn**—Virtual private network



NOTE: EX Series switches run a subset of these protocols. See the switch CLI for details.

Required Privilege Level	view
List of Sample Output	show route protocol access on page 908 show route protocol access-internal extensive on page 908 show route protocol arp on page 908 show route protocol bgp on page 909 show route protocol bgp detail on page 909 show route protocol bgp extensive on page 909 show route protocol bgp terse on page 910 show route protocol direct on page 910 show route protocol frr on page 911 show route protocol l2circuit detail on page 911 show route protocol l2vpn extensive on page 912 show route protocol ldp on page 913 show route protocol ldp extensive on page 913 show route protocol ospf (Layer 3 VPN) on page 914 show route protocol ospf detail on page 915 show route protocol rip on page 915 show route protocol rip detail on page 915 show route protocol ripng table inet6 on page 916 show route protocol static detail on page 916

Output Fields For information about output fields, see the output field tables for the [show route](#) command, the [show route detail](#) command, the [show route extensive](#) command, or the [show route terse](#) command.

Sample Output

show route protocol access

```
user@host> show route protocol access
inet.0: 30380 destinations, 30382 routes (30379 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

13.160.0.3/32      *[Access/13] 00:00:09
                  > to 13.160.0.2 via fe-0/0/0.0
13.160.0.4/32      *[Access/13] 00:00:09
                  > to 13.160.0.2 via fe-0/0/0.0
13.160.0.5/32      *[Access/13] 00:00:09
                  > to 13.160.0.2 via fe-0/0/0.0
```

show route protocol access-internal extensive

```
user@host> show route protocol access-internal 13.160.0.19 extensive
inet.0: 100020 destinations, 100022 routes (100019 active, 0 holddown, 1 hidden)
13.160.0.19/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 13.160.0.19/32 -> {13.160.0.2}
    *Access-internal Preference: 12
        Next-hop reference count: 200000
        Next hop: 13.160.0.2 via fe-0/0/0.0, selected
        State: <Active Int>
    Age: 36
        Task: RPD Unix Domain Server./var/run/rpd_serv.local
        Announcement bits (1): 0-KRT
        AS path: I
```

show route protocol arp

```
user@host> show route protocol arp
inet.0: 43 destinations, 43 routes (42 active, 0 holddown, 1 hidden)

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

cust1.inet.0: 1033 destinations, 2043 routes (1033 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.1.3/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.4/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.5/32      [ARP/4294967293] 00:04:32, from 20.20.1.1
                  Unusable
20.20.1.6/32      [ARP/4294967293] 00:04:34, from 20.20.1.1
                  Unusable
20.20.1.7/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.8/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.9/32      [ARP/4294967293] 00:04:35, from 20.20.1.1
                  Unusable
20.20.1.10/32     [ARP/4294967293] 00:04:35, from 20.20.1.1
```

```

Unusable
20.20.1.11/32      [ARP/4294967293] 00:04:33, from 20.20.1.1
Unusable
20.20.1.12/32      [ARP/4294967293] 00:04:33, from 20.20.1.1
Unusable
20.20.1.13/32      [ARP/4294967293] 00:04:33, from 20.20.1.1
Unusable
...

```

show route protocol bgp

```

user@host> show route protocol bgp 192.168.64.0/21
inet.0: 335832 destinations, 335833 routes (335383 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.64.0/21      *[BGP/170] 6d 10:41:16, localpref 100, from 192.168.69.71
AS path: 10458 14203 2914 4788 4788 I
> to 192.168.167.254 via fxp0.0

```

show route protocol bgp detail

```

user@host> show route protocol bgp 66.117.63.0/24 detail
inet.0: 335805 destinations, 335806 routes (335356 active, 0 holddown, 450 hidden)
66.117.63.0/24      (1 entry, 1 announced)
  *BGP      Preference: 170/-101
              Next hop type: Indirect
              Next-hop reference count: 1006436
              Source: 192.168.69.71
              Next hop type: Router, Next hop index: 324
              Next hop: 192.168.167.254 via fxp0.0, selected
              Protocol next hop: 192.168.69.71
              Indirect next hop: 8e166c0 342
              State: <Active Ext>
              Local AS: 69 Peer AS: 10458
              Age: 6d 10:42:42 Metric2: 0
              Task: BGP_10458.192.168.69.71+179
              Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree

1
  AS path: 10458 14203 2914 4788 4788 I
  Communities: 2914:410 2914:2403 2914:3400
  Accepted
  Localpref: 100
  Router ID: 207.17.136.192

```

show route protocol bgp extensive

```

user@host> show route protocol bgp 192.168.64.0/21 extensive

inet.0: 335827 destinations, 335828 routes (335378 active, 0 holddown, 450 hidden)
192.168.64.0/21 (1 entry, 1 announced)
TSI:
KRT in-kernel 1.9.0.0/16 -> {indirect(342)}
Page 0 idx 1 Type 1 val db31a80
  Nexthop: Self
  AS path: [69] 10458 14203 2914 4788 4788 I
  Communities: 2914:410 2914:2403 2914:3400
Path 1.9.0.0 from 192.168.69.71 Vector len 4. Val: 1
  *BGP      Preference: 170/-101
              Next hop type: Indirect
              Next-hop reference count: 1006502
              Source: 192.168.69.71
              Next hop type: Router, Next hop index: 324

```

```

Next hop: 192.168.167.254 via fxp0.0, selected
Protocol next hop: 192.168.69.71
Indirect next hop: 8e166c0 342
State: <Active Ext>
Local AS: 69 Peer AS: 10458
Age: 6d 10:44:45 Metric2: 0
Task: BGP_10458.192.168.69.71+179
Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree
1
AS path: 10458 14203 2914 4788 4788 I
Communities: 2914:410 2914:2403 2914:3400
Accepted
Localpref: 100
Router ID: 207.17.136.192
Indirect next hops: 1
  Protocol next hop: 192.168.69.71
  Indirect next hop: 8e166c0 342
  Indirect path forwarding next hops: 1
    Next hop type: Router
    Next hop: 192.168.167.254 via fxp0.0
  192.168.0.0/16 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
    Nexthop: 192.168.167.254 via fxp0.0

```

show route protocol bgp terse

```

user@host> show route protocol bgp 192.168.64.0/21 terse

inet.0: 24 destinations, 32 routes (23 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
192.168.64.0/21    B 170      100          >100.1.3.2    10023 21 I

```

show route protocol direct

```

user@host> show route protocol direct

inet.0: 335843 destinations, 335844 routes (335394 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both

8.8.8.0/24          *[Direct/0] 17w0d 10:31:49
> via fe-1/3/1.0
10.255.165.1/32     *[Direct/0] 25w4d 04:13:18
> via lo0.0
30.30.30.0/24       *[Direct/0] 17w0d 23:06:26
> via fe-1/3/2.0
192.168.164.0/22    *[Direct/0] 25w4d 04:13:20
> via fxp0.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

47.0005.80ff.f800.0000.0108.0001.0102.5516.5001/152
*[Direct/0] 25w4d 04:13:21
> via lo0.0

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```



```

abcd::10:255:165:1/128
    *[Direct/0] 25w4d 04:13:21
    > via lo0.0
fe80::2a0:a5ff:fe12:ad7/128
    *[Direct/0] 25w4d 04:13:21
    > via lo0.0

```

show route protocol frr

```

user@host> show route protocol frr
inet.0: 43 destinations, 43 routes (42 active, 0 holddown, 1 hidden)

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

cust1.inet.0: 1033 destinations, 2043 routes (1033 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.1.3/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.3 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.4/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.4 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.5/32      *[FRR/200] 00:05:35, from 20.20.1.1
                  > to 20.20.1.5 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.6/32      *[FRR/200] 00:05:37, from 20.20.1.1
                  > to 20.20.1.6 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.7/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.7 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.8/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.8 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.9/32      *[FRR/200] 00:05:38, from 20.20.1.1
                  > to 20.20.1.9 via ge-4/1/0.0
                  to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.10/32     *[FRR/200] 00:05:38, from 20.20.1.1
...

```

show route protocol l2circuit detail

```

user@host> show route protocol l2circuit detail

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
100000 (1 entry, 1 announced)
    *L2CKT Preference: 7
        Next hop: via ge-2/0/0.0, selected
        Label operation: Pop      Offset: 4
        State: <Active Int>
        Local AS: 99
        Age: 9:52
        Task: Common L2 VC
        Announcement bits (1): 0-KRT
        AS path: I

ge-2/0/0.0 (1 entry, 1 announced)
    *L2CKT Preference: 7
        Next hop: via so-1/1/2.0 weight 1, selected
        Label-switched-path my-lsp

```

```

Label operation: Push 100000, Push 100000(top)[0] Offset: -4
Protocol next hop: 10.245.255.63
Push 100000 Offset: -4
  Indirect next hop: 86af0c0 298
State: <Active Int>
Local AS: 99
Age: 9:52
Task: Common L2 VC
Announcement bits (2): 0-KRT 1-Common L2 VC
AS path: I

```

```
l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
```

```

10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
  *L2CKT Preference: 7
    Next hop: via so-1/1/2.0 weight 1, selected
    Label-switched-path my-lsp
    Label operation: Push 100000[0]
    Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
    State: <Active Int>
    Local AS: 99
    Age: 10:21
    Task: l2 circuit
    Announcement bits (1): 0-LDP
    AS path: I
    VC Label 100000, MTU 1500, VLAN ID 512

```

show route protocol l2vpn extensive

```
user@host> show route protocol l2vpn extensive
```

```
inet.0: 14 destinations, 15 routes (13 active, 0 holddown, 1 hidden)
```

```
inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

```
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

```

mpls.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
800001 (1 entry, 1 announced)

```

```
TSI:
```

```
KRT in-kernel 800001 /36 -> {so-0/0/0.0}
```

```

  *L2VPN Preference: 7
    Next hop: via so-0/0/0.0 weight 49087 balance 97%, selected
    Label operation: Pop      Offset: 4
    State: <Active Int>
    Local AS: 69
    Age: 7:48
    Task: Common L2 VC
    Announcement bits (1): 0-KRT
    AS path: I

```

```
so-0/0/0.0 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel so-0/0/0.0 /16 -> {indirect(288)}
```

```

  *L2VPN Preference: 7
    Next hop: via so-0/0/1.0, selected
    Label operation: Push 800000 Offset: -4
    Protocol next hop: 10.255.14.220
    Push 800000 Offset: -4
    Indirect next hop: 85142a0 288
    State: <Active Int>

```

```

Local AS:    69
Age: 7:48
Task: Common L2 VC
Announcement bits (2): 0-KRT 1-Common L2 VC
AS path: I
Communities: target:69:1 Layer2-info: encaps:PPP,
control flags:2, mtu: 0

```

show route protocol ldp

```

user@host> show route protocol ldp
inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.16.1/32    *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Push 100000
192.168.17.1/32    *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0

private1___.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100064            *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Pop
100064(S=0)        *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Pop
100080            *[LDP/9] 1d 23:03:35, metric 1
                  > via t1-4/0/0.0, Swap 100000

```

show route protocol ldp extensive

```

user@host> show route protocol ldp extensive
192.168.16.1/32 (1 entry, 1 announced)
  State: <FlashAll>
  *LDP    Preference: 9
          Next-hop reference count: 3
          Next hop: via t1-4/0/0.0, selected
          Label operation: Push 100000
          State: <Active Int>
          Local AS: 65500
          Age: 1d 23:03:58      Metric: 1
          Task: LDP
          Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
          AS path: I

192.168.17.1/32 (1 entry, 1 announced)
  State: <FlashAll>
  *LDP    Preference: 9
          Next-hop reference count: 3
          Next hop: via t1-4/0/0.0, selected
          State: <Active Int>
          Local AS: 65500
          Age: 1d 23:03:58      Metric: 1
          Task: LDP
          Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
          AS path: I

```

```
private1__inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
```

```
mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
```

```
100064 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 100064 /36 -> {t1-4/0/0.0}
```

```
*LDP      Preference: 9
           Next-hop reference count: 2
           Next hop: via t1-4/0/0.0, selected
           State: <Active Int>
           Local AS: 65500
           Age: 1d 23:03:58      Metric: 1
           Task: LDP
           Announcement bits (1): 0-KRT
           AS path: I
           Prefixes bound to route: 192.168.17.1/32
```

```
100064(S=0) (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 100064 /40 -> {t1-4/0/0.0}
```

```
*LDP      Preference: 9
           Next-hop reference count: 2
           Next hop: via t1-4/0/0.0, selected
           Label operation: Pop
           State: <Active Int>
           Local AS: 65500
           Age: 1d 23:03:58      Metric: 1
           Task: LDP
           Announcement bits (1): 0-KRT
           AS path: I
```

```
100080 (1 entry, 1 announced)
```

```
TSI:
```

```
KRT in-kernel 100080 /36 -> {t1-4/0/0.0}
```

```
*LDP      Preference: 9
           Next-hop reference count: 2
           Next hop: via t1-4/0/0.0, selected
           Label operation: Swap 100000
           State: <Active Int>
           Local AS: 65500
           Age: 1d 23:03:58      Metric: 1
           Task: LDP
           Announcement bits (1): 0-KRT
           AS path: I
           Prefixes bound to route: 192.168.16.1/32
```

show route protocol ospf (Layer 3 VPN)

```
user@host> show route protocol ospf
```

```
inet.0: 40 destinations, 40 routes (39 active, 0 holddown, 1 hidden)
```

```
+ = Active Route, - = Last Active, * = Both
```

```
10.39.1.4/30      *[OSPF/10] 00:05:18, metric 4
                  > via t3-3/2/0.0
10.39.1.8/30      [OSPF/10] 00:05:18, metric 2
                  > via t3-3/2/0.0
10.255.14.171/32  *[OSPF/10] 00:05:18, metric 4
                  > via t3-3/2/0.0
10.255.14.179/32 *[OSPF/10] 00:05:18, metric 2
                  > via t3-3/2/0.0
```

```

224.0.0.5/32      *[OSPF/10] 20:25:55, metric 1

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30     [OSPF/10] 00:05:43, metric 1
                  > via so-0/2/2.0
10.255.14.173/32  *[OSPF/10] 00:05:43, metric 1
                  > via so-0/2/2.0
224.0.0.5/32      *[OSPF/10] 20:26:20, metric 1

```

show route protocol ospf detail

```

user@host> show route protocol ospf detail
VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30 (2 entries, 0 announced)
    OSPF    Preference: 10
            Nexthop: via so-0/2/2.0, selected
            State: <Int>
            Inactive reason: Route Preference
            Age: 6:25      Metric: 1
            Area: 0.0.0.0
            Task: VPN-AB-OSPF
            AS path: I
            Communities: Route-Type:0.0.0.0:1:0

...

```

show route protocol rip

```

user@host> show route protocol rip
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32  *[RIP/100] 20:24:34, metric 2
                  > to 10.39.1.22 via t3-0/2/2.0
224.0.0.9/32      *[RIP/100] 00:03:59, metric 1

```

show route protocol rip detail

```

user@host> show route protocol rip detail
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32 (1 entry, 1 announced)
    *RIP    Preference: 100
            Nexthop: 10.39.1.22 via t3-0/2/2.0, selected
            State: <Active Int>
            Age: 20:25:02  Metric: 2
            Task: VPN-AB-RIPv2
            Announcement bits (2): 0-KRT 2-BGP.0.0.0.0+179
            AS path: I
            Route learned from 10.39.1.22 expires in 96 seconds

```

show route protocol ripng table inet6

```
user@host> show route protocol ripng table inet6
inet6.0: 4215 destinations, 4215 routes (4214 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

1111::1/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::2/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::3/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::4/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::5/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::6/128      *[RIPng/100] 02:13:33, metric 2
                  > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
```

show route protocol static detail

```
user@host> show route protocol static detail
inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
10.5.0.0/16 (1 entry, 1 announced)
    *Static Preference: 5
        Next hop type: Router, Next hop index: 324
        Address: 0x9274010
        Next-hop reference count: 27
        Next hop: 192.168.187.126 via fxp0.0, selected
        Session Id: 0x0
        State: <Active NoReadvrt Int Ext>
        Age: 7w3d 21:24:25
        Validation State: unverified
        Task: RT
        Announcement bits (1): 0-KRT
        AS path: I

10.10.0.0/16 (1 entry, 1 announced)
    *Static Preference: 5
        Next hop type: Router, Next hop index: 324
        Address: 0x9274010
        Next-hop reference count: 27
        Next hop: 192.168.187.126 via fxp0.0, selected
        Session Id: 0x0
        State: <Active NoReadvrt Int Ext>
        Age: 7w3d 21:24:25
        Validation State: unverified
        Task: RT
        Announcement bits (1): 0-KRT
        AS path: I

10.13.10.0/23 (1 entry, 1 announced)
    *Static Preference: 5
        Next hop type: Router, Next hop index: 324
        Address: 0x9274010
        Next-hop reference count: 27
        Next hop: 192.168.187.126 via fxp0.0, selected
        Session Id: 0x0
        State: <Active NoReadvrt Int Ext>
        Age: 7w3d 21:24:25
        Validation State: unverified
```

Task: RT
Announcement bits (1): 0-KRT
AS path: I

show route range

Syntax	<code>show route range</code> <code><brief detail extensive terse></code> <code><destination-prefix></code> <code><logical-system (all <i>logical-system-name</i>)></code>
Syntax (EX Series Switches)	<code>show route range</code> <code><brief detail extensive terse></code> <code><destination-prefix></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display routing table entries using a prefix range.
Options	<p>none—Display standard information about all routing table entries using a prefix range.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>destination-prefix—Destination and prefix mask for the range.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show route range on page 918 show route range destination-prefix on page 919 show route range detail on page 919 show route range extensive on page 920 show route range terse on page 921
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route range

```
user@host> show route range

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16      *[Static/5] 00:30:01
                  > to 192.168.71.254 via fxp0.0
10.209.0.0/16    *[Static/5] 00:30:01
                  > to 192.168.71.254 via fxp0.0
10.255.71.14/32 *[Direct/0] 00:30:01
                  > via lo0.0
```



```

172.16.0.0/12      *[Static/5] 00:30:01
                  > to 192.168.71.254 via fxp0.0
192.168.0.0/16    *[Static/5] 00:30:01
                  > to 192.168.71.254 via fxp0.0
192.168.64.0/21   *[Direct/0] 00:30:01
                  > via fxp0.0
192.168.71.14/32  *[Local/0] 00:30:01
                  Local via fxp0.0
192.168.102.0/23  *[Static/5] 00:30:01
                  > to 192.168.71.254 via fxp0.0
...

```

show route range destination-prefix

```

user@host> show route range 192.168.0.0/16

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.0.0/16    *[Static/5] 00:31:14
                  > to 192.168.71.254 via fxp0.0
192.168.64.0/21   *[Direct/0] 00:31:14
                  > via fxp0.0
192.168.71.14/32  *[Local/0] 00:31:14
                  Local via fxp0.0
192.168.102.0/23  *[Static/5] 00:31:14
                  > to 192.168.71.254 via fxp0.0

```

show route range detail

```

user@host> show route range detail

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:05
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.209.0.0/16 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:05
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.255.71.14/32 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Active Int>
    Age: 30:05
    Task: IF

```

```
AS path: I
172.16.0.0/12 (1 entry, 1 announced)
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:05
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I
```

...

show route range extensive

```
user@host> show route range extensive

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:17
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.209.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:17
    Task: RT
    Announcement bits (1): 0-KRT
    AS path: I

10.255.71.14/32 (1 entry, 0 announced)
  *Direct Preference: 0
    Next hop type: Interface
    Next-hop reference count: 1
    Next hop: via lo0.0, selected
    State: <Active Int>
    Age: 30:17
    Task: IF
    AS path: I

172.16.0.0/12 (1 entry, 1 announced)
TSI:
KRT in-kernel 172.16.0.0/12 -> {192.168.71.254}
  *Static Preference: 5
    Next-hop reference count: 22
    Next hop: 192.168.71.254 via fxp0.0, selected
    State: <Active NoReadvrt Int Ext>
    Age: 30:17
```

Task: RT
 Announcement bits (1): 0-KRT
 AS path: I

...

show route range terse

user@host> show route range terse

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
 + = Active Route, - = Last Active, * = Both

A	Destination	P	Prf	Metric 1	Metric 2	Next hop	AS path
*	10.10.0.0/16	S	5			>192.168.71.254	
*	10.209.0.0/16	S	5			>192.168.71.254	
*	10.255.71.14/32	D	0			>lo0.0	
*	172.16.0.0/12	S	5			>192.168.71.254	
*	192.168.0.0/16	S	5			>192.168.71.254	
*	192.168.64.0/21	D	0			>fxp0.0	
*	192.168.71.14/32	L	0			Local	
*	192.168.102.0/23	S	5			>192.168.71.254	
*	207.17.136.0/24	S	5			>192.168.71.254	
*	207.17.136.192/32	S	5			>192.168.71.254	

__juniper_private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both

A	Destination	P	Prf	Metric 1	Metric 2	Next hop	AS path
*	10.0.0.0/8	D	0			>fxp2.0	
		D	0			>fxp1.0	
*	10.0.0.4/32	L	0			Local	

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both

A	Destination	P	Prf	Metric 1	Metric 2	Next hop	AS path
	47.0005.80ff.f800.0000.0108.0001.0102.5507.1014/152						
*		D	0			>lo0.0	

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both

A	Destination	P	Prf	Metric 1	Metric 2	Next hop	AS path
	abcd::10:255:71:14/128						
*		D	0			>lo0.0	
	fe80::280:42ff:fe11:226f/128						
*		D	0			>lo0.0	

__juniper_private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both

A	Destination	P	Prf	Metric 1	Metric 2	Next hop	AS path
	fe80::280:42ff:fe11:226f/128						
*		D	0			>lo0.16385	

show route receive-protocol

Syntax	show route receive-protocol <i>protocol neighbor-address</i> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>	
Syntax (EX Series Switches)	show route receive-protocol <i>protocol neighbor-address</i> <brief detail extensive terse>	
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.	
Description	Display the routing information as it was received through a particular neighbor using a particular dynamic routing protocol.	
Options	brief detail extensive terse —(Optional) Display the specified level of output. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. <i>protocol neighbor-address</i> —Protocol transmitting the route (bgp , dvmrp , msdp , pim , rip , or ripng) and address of the neighboring router from which the route entry was received.	
Additional Information	The output displays the selected routes and the attributes with which they were received, but does not show the effects of import policy on the routing attributes.	
Required Privilege Level	view	
List of Sample Output	show route receive-protocol bgp on page 925 show route receive-protocol bgp extensive on page 925 show route receive-protocol bgp table extensive on page 925 show route receive-protocol bgp logical-system extensive on page 926 show route receive-protocol bgp detail (Layer 2 VPN) on page 927 show route receive-protocol bgp extensive (Layer 2 VPN) on page 927 show route receive-protocol bgp (Layer 3 VPN) on page 928 show route receive-protocol bgp detail (Layer 3 VPN) on page 928 show route receive-protocol bgp extensive (Layer 3 VPN) on page 929	
Output Fields	Table 139 on page 922 describes the output fields for the show route receive-protocol command. Output fields are listed in the approximate order in which they appear.	

Table 139: show route receive-protocol Output Fields

Field Name	Field Description	Level of Output
<i>routing-table-name</i>	Name of the routing table—for example, inet.0.	All levels
<i>number destinations</i>	Number of destinations for which there are routes in the routing table.	All levels

Table 139: show route receive-protocol Output Fields (*continued*)

Field Name	Field Description	Level of Output
<i>number routes</i>	Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> • active • holddown (routes that are in pending state before being declared inactive) • hidden (routes that are not used because of a routing policy) 	All levels
Prefix	Destination prefix.	none brief
MED	Multiple exit discriminator value included in the route.	none brief
<i>destination-prefix</i> (entry, announced)	Destination prefix. The entry value is the number of routes for this destination, and the announced value is the number of routes being announced for this destination.	detail extensive
Route Distinguisher	64-bit prefix added to IP subnets to make them unique.	detail extensive
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.	detail extensive
VPN Label	Virtual private network (VPN) label. Packets are sent between CE and PE routing devices by advertising VPN labels. VPN labels transit over either an RSVP or an LDP label-switched path (LSP) tunnel.	detail extensive
Next hop	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.	All levels
Localpref or Lclpref	Local preference value included in the route.	All levels

Table 139: show route receive-protocol Output Fields (*continued*)

Field Name	Field Description	Level of Output
AS path	<p>Autonomous system (AS) path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • ?—Incomplete; typically, the AS path was aggregated. <p>When AS path numbers are included in the route, the format is as follows:</p> <ul style="list-style-type: none"> • []—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used the AS-path merge process, as defined in RFC 4893. • []—If more than one AS number is configured on the router, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path. • { }—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • ()—Parentheses enclose a confederation. • ([])—Parentheses and brackets enclose a confederation set. <p>NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.</p>	All levels
Cluster list	(For route reflected output only) Cluster ID sent by the route reflector.	detail extensive
Originator ID	(For route reflected output only) Address of routing device that originally sent the route to the route reflector.	detail extensive
Communities	Community path attribute for the route. See the Output Field table in the show route detail command for all possible values for this field.	detail extensive
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.	detail extensive
Attrset AS	Number, local preference, and path of the AS that originated the route. These values are stored in the Attrset attribute at the originating routing device.	detail extensive
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).	detail extensive
control flags	Control flags: none or Site Down .	detail extensive
mtu	Maximum transmission unit (MTU) of the Layer 2 circuit.	detail extensive

Sample Output

show route receive-protocol bgp

```
user@host> show route receive-protocol bgp 10.255.245.215

inet.0: 28 destinations, 33 routes (27 active, 0 holddown, 1 hidden)
Prefix          Next hop          MED      Lclpref  AS path
10.22.1.0/24     10.255.245.215    0        100      I
10.22.2.0/24     10.255.245.215    0        100      I
```

show route receive-protocol bgp extensive

```
user@host> show route receive-protocol bgp 10.255.245.63 extensive
inet.0: 244 destinations, 244 routes (243 active, 0 holddown, 1 hidden)
Prefix          Next hop          MED      Lclpref  AS path
1.1.1.0/24 (1 entry, 1 announced)
  Next hop: 10.0.50.3
  Localpref: 100
  AS path: I <Originator>
  Cluster list: 10.2.3.1
  Originator ID: 10.255.245.45
165.3.0.0/16 (1 entry, 1 announced)
  Next hop: 111.222.5.254
  Localpref: 100
  AS path: I <Originator>
  Cluster list: 10.2.3.1
  Originator ID: 10.255.245.68
165.4.0.0/16 (1 entry, 1 announced)
  Next hop: 111.222.5.254
  Localpref: 100
  AS path: I <Originator>
  Cluster list: 10.2.3.1
  Originator ID: 10.255.245.45
195.1.2.0/24 (1 entry, 1 announced)
  Next hop: 111.222.5.254
  Localpref: 100
  AS path: I <Originator>
  Cluster list: 10.2.3.1
  Originator ID: 10.255.245.68
inet.2: 63 destinations, 63 routes (63 active, 0 holddown, 0 hidden)
Prefix          Next hop          MED      Lclpref  AS path
inet.3: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix          Next hop          MED      Lclpref  AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Next hop          MED      Lclpref  AS path
mpls.0: 48 destinations, 48 routes (48 active, 0 holddown, 0 hidden)
```

show route receive-protocol bgp table extensive

```
user@host> show route receive-protocol bgp 207.17.136.192 table inet.0 66.117.68.0/24 extensive
inet.0: 227315 destinations, 227316 routes (227302 active, 0 holddown, 13 hidden)
* 66.117.63.0/24 (1 entry, 1 announced)
  Nexthop: 207.17.136.29
  Localpref: 100
  AS path: AS2 PA[6]: 14203 2914 3356 29748 33437 AS_TRANS
  AS path: AS4 PA[2]: 33437 393219
  AS path: Merged[6]: 14203 2914 3356 29748 33437 393219 I
  Communities: 2914:420
```

show route receive-protocol bgp 10.0.0.9 logical-system PE4 extensive

```
user@host> show route receive-protocol bgp 10.0.0.9 logical-system PE4 extensive
inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)
* 10.0.0.0/30 (1 entry, 1 announced)
  Accepted
  Route Label: 3
  Nexthop: 10.0.0.9
  AS path: 13979 I

* 10.0.0.4/30 (1 entry, 1 announced)
  Accepted
  Route Label: 3
  Nexthop: 10.0.0.9
  AS path: 13979 I

10.0.0.8/30 (2 entries, 1 announced)
  Accepted
  Route Label: 3
  Nexthop: 10.0.0.9
  AS path: 13979 I

* 10.9.9.1/32 (1 entry, 1 announced)
  Accepted
  Route Label: 3
  Nexthop: 10.0.0.9
  AS path: 13979 I

* 10.100.1.1/32 (1 entry, 1 announced)
  Accepted
  Route Label: 3
  Nexthop: 10.0.0.9
  AS path: 13979 I

* 44.0.0.0/24 (1 entry, 1 announced)
  Accepted
  Route Label: 300096
  Nexthop: 10.0.0.9
  AS path: 13979 I
  AIGP: 203

* 55.0.0.0/24 (1 entry, 1 announced)
  Accepted
  Route Label: 300112
  Nexthop: 10.0.0.9
  AS path: 13979 7018 I
  AIGP: 25

* 66.0.0.0/24 (1 entry, 1 announced)
  Accepted
  Route Label: 300144
  Nexthop: 10.0.0.9
  AS path: 13979 7018 I

* 99.0.0.0/24 (1 entry, 1 announced)
  Accepted
  Route Label: 300160
  Nexthop: 10.0.0.9
  AS path: 13979 7018 I
```


show route receive-protocol bgp detail (Layer 2 VPN)

```

user@host> show route receive-protocol bgp 10.255.14.171 detail
inet.0: 68 destinations, 68 routes (67 active, 0 holddown, 1 hidden)
Prefix          Nexthop          MED      Lclpref AS path
inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
frame-vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0
hidden)
Prefix          Nexthop          MED      Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 1 announced)
  Route Distinguisher: 10.255.245.35:1
  Label-base : 800000, range : 4, status-vector : 0x0
  Nexthop: 10.255.245.35
  Localpref: 100
  AS path: I
  Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
control flags: 0, mtu: 0
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 0 announced)
  Route Distinguisher: 10.255.245.35:1
  Label-base : 800000, range : 4, status-vector : 0x0
  Nexthop: 10.255.245.35
  Localpref: 100
  AS path: I
  Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
control flags:0, mtu: 0

```

show route receive-protocol bgp extensive (Layer 2 VPN)

```

user@host> show route receive-protocol bgp 10.255.14.171 extensive
inet.0: 68 destinations, 68 routes (67 active, 0 holddown, 1 hidden)
Prefix          Nexthop          MED      Lclpref AS path
inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
frame-vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 1 announced)
  Route Distinguisher: 10.255.245.35:1
  Label-base : 800000, range : 4, status-vector : 0x0
  Nexthop: 10.255.245.35
  Localpref: 100
  AS path: I
  Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
control flags:0, mtu: 0
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED      Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 0 announced)
  Route Distinguisher: 10.255.245.35:1
  Label-base : 800000, range : 4, status-vector : 0x0
  Nexthop: 10.255.245.35
  Localpref: 100

```

```

AS path: I
Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
control flags:0, mtu: 0

```

show route receive-protocol bgp (Layer 3 VPN)

```

user@host> show route receive-protocol bgp 10.255.14.171
inet.0: 33 destinations, 33 routes (32 active, 0 holddown, 1 hidden)
Prefix          Nexthop          MED    Lclpref AS path
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
VPN-A.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
10.255.14.175/32 10.255.14.171          100 2 I
10.255.14.179/32 10.255.14.171          2    100 I
VPN-B.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
10.255.14.175/32 10.255.14.171          100 2 I
10.255.14.177/32 10.255.14.171          100 I
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
mpls.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
bgp.l3vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Prefix          Nexthop          MED    Lclpref AS path
10.255.14.171:300:10.255.14.177/32
                  10.255.14.171          100 I
10.255.14.171:100:10.255.14.179/32
                  10.255.14.171          2    100 I
10.255.14.171:200:10.255.14.175/32
                  10.255.14.171          100 2 I

```

show route receive-protocol bgp detail (Layer 3 VPN)

```

user@host> show route receive-protocol bgp 10.255.14.174 detail
inet.0: 16 destinations, 17 routes (15 active, 0 holddown, 1 hidden)
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
vpna.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
* 10.49.0.0/30 (1 entry, 1 announced)
  Route Distinguisher: 10.255.14.176:2
  VPN Label: 101264
  Nexthop: 10.255.14.174
  Localpref: 100
  AS path: I
  Communities: target:200:100
  AttrSet AS: 100
    Localpref: 100
    AS path: I
* 10.255.14.172/32 (1 entry, 1 announced)
  Route Distinguisher: 10.255.14.176:2
  VPN Label: 101280
  Nexthop: 10.255.14.174
  Localpref: 100
  AS path: I
  Communities: target:200:100
  AttrSet AS: 100
    Localpref: 100
    AS path: I
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

```

```

* 10.255.14.174:2:10.49.0.0/30 (1 entry, 0 announced)
  Route Distinguisher: 10.255.14.174:2
  VPN Label: 101264
  Nexthop: 10.255.14.174
  Localpref: 100
  AS path: I
  Communities: target:200:100
  AttrSet AS: 100
    Localpref: 100
    AS path: I
* 10.255.14.174:2:10.255.14.172/32 (1 entry, 0 announced)
  Route Distinguisher: 10.255.14.174:2
  VPN Label: 101280
  Nexthop: 10.255.14.174
  Localpref: 100
  AS path: I
  Communities: target:200:100
  AttrSet AS: 100
    Localpref: 100
    AS path: I
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

```

show route receive-protocol bgp extensive (Layer 3 VPN)

```

user@host> show route receive-protocol bgp 10.255.245.63 extensive
inet.0: 244 destinations, 244 routes (243 active, 0 holddown, 1 hidden)
  Prefix          Nexthop          MED      Lclpref AS path
  1.1.1.0/24 (1 entry, 1 announced)
    Nexthop: 10.0.50.3
    Localpref: 100
    AS path: I <Originator>
    Cluster list: 10.2.3.1
    Originator ID: 10.255.245.45
  165.3.0.0/16 (1 entry, 1 announced)
    Nexthop: 111.222.5.254
    Localpref: 100
    AS path: I <Originator>
    Cluster list: 10.2.3.1
    Originator ID: 10.255.245.68
  165.4.0.0/16 (1 entry, 1 announced)
    Nexthop: 111.222.5.254
    Localpref: 100
    AS path: I <Originator>
    Cluster list: 10.2.3.1
    Originator ID: 10.255.245.45
  195.1.2.0/24 (1 entry, 1 announced)
    Nexthop: 111.222.5.254
    Localpref: 100
    AS path: I <Originator>
    Cluster list: 10.2.3.1
    Originator ID: 10.255.245.68
inet.2: 63 destinations, 63 routes (63 active, 0 holddown, 0 hidden)
  Prefix          Nexthop          MED      Lclpref AS path
inet.3: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
  Prefix          Nexthop          MED      Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
  Prefix          Nexthop          MED      Lclpref AS path
mpls.0: 48 destinations, 48 routes (48 active, 0 holddown, 0 hidden)

```

show route resolution

Syntax	<code>show route resolution</code> <code><brief detail extensive summary></code> <code><index <i>index</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code> <code><prefix></code> <code><table <i>routing-table-name</i>></code> <code><unresolved></code>
Syntax (EX Series Switches)	<code>show route resolution</code> <code><brief detail extensive summary></code> <code><index <i>index</i>></code> <code><prefix></code> <code><table <i>routing-table-name</i>></code> <code><unresolved></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the entries in the next-hop resolution database. This database provides for recursive resolution of next hops through other prefixes in the routing table.
Options	none —Display standard information about all entries in the next-hop resolution database. brief detail extensive summary —(Optional) Display the specified level of output. index <i>index</i> —(Optional) Show the index of the resolution tree. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. prefix <i>network/destination-prefix</i> —(Optional) Display database entries for the specified address. table <i>routing-table-name</i> —(Optional) Display information about a particular routing table (for example, <i>inet.0</i>) where policy-based export is currently enabled. unresolved —(Optional) Display routes that could not be resolved.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• <i>Example: Configuring Route Resolution on PE Routers</i>
List of Sample Output	show route resolution detail on page 931 show route resolution summary on page 932 show route resolution unresolved on page 932

Output Fields Table 140 on page 931 describes the output fields for the **show route resolution** command. Output fields are listed in the approximate order in which they appear.

Table 140: show route resolution Output Fields

Field Name	Field Description
<i>routing-table-name</i>	Name of the routing table whose prefixes are resolved using the entries in the route resolution database. For routing table groups, this is the name of the primary routing table whose prefixes are resolved using the entries in the route resolution database.
Tree index	Tree index identifier.
Nodes	Number of nodes in the tree.
Reference count	Number of references made to the next hop.
Contributing routing tables	Routing tables used for next-hop resolution.
Originating RIB	Name of the routing table whose active route was used to determine the forwarding next-hop entry in the resolution database. For example, in the case of inet.0 resolving through inet.0 and inet.3 , this field indicates which routing table, inet.0 or inet.3 , provided the best path for a particular prefix.
Metric	Metric associated with the forwarding next hop.
Node path count	Number of nodes in the path.
Forwarding next hops	Number of forwarding next hops. The forwarding next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.

Sample Output

show route resolution detail

```

user@host> show route resolution detail
Tree Index: 1, Nodes 0, Reference Count 1
Contributing routing tables: inet.3
Tree Index: 2, Nodes 23, Reference Count 1
Contributing routing tables: inet.0 inet.3
10.10.0.0/16 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
10.31.1.0/30 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
10.31.1.1/32 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 0
10.31.1.4/30 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
10.31.1.5/32 Originating RIB: inet.0

```

```
Node path count: 1
Forwarding nexthops: 0
10.31.2.0/30 Originating RIB: inet.0
Metric: 2 Node path count: 1
Forwarding nexthops: 2
10.31.11.0/24 Originating RIB: inet.0
Node path count: 1
Forwarding nexthops: 1
```

show route resolution summary

```
user@host> show route resolution summary
Tree Index: 1, Nodes 24, Reference Count 1
Contributing routing tables: :voice.inet.0 :voice.inet.3
Tree Index: 2, Nodes 2, Reference Count 1
Contributing routing tables: inet.3
Tree Index: 3, Nodes 43, Reference Count 1
Contributing routing tables: inet.0 inet.3
```

show route resolution unresolved

```
user@host> show route resolution unresolved
Tree Index 1
vt-3/2/0.32769.0      /16
    Protocol Nexthop: 10.255.71.238 Push 800000
    Indirect nexthop: 0 -
vt-3/2/0.32772.0      /16
    Protocol Nexthop: 10.255.70.103 Push 800008
    Indirect nexthop: 0 -
Tree Index 2
```

show route snooping

Syntax	<pre>show route snooping <brief detail extensive terse> <all> <best address/prefix> <exact address> <range prefix-range> <summary> <table table-name></pre>
Release Information	<p>Command introduced in Junos OS Release 8.5.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	Display the entries in the routing table that were learned from snooping.
Options	<p>none—Display the entries in the routing table that were learned from snooping.</p> <p>brief detail extensive terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.</p> <p>all—(Optional) Display all entries, including hidden entries.</p> <p>best address/prefix—(Optional) Display the longest match for the provided address and optional prefix.</p> <p>exact address/prefix—(Optional) Display exact matches for the provided address and optional prefix.</p> <p>range prefix-range—(Optional) Display information for the provided address range.</p> <p>summary—(Optional) Display route snooping summary statistics.</p> <p>table table-name—(Optional) Display information for the named table.</p>
Required Privilege Level	view
List of Sample Output	show route snooping detail on page 933
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route snooping detail

```
user@host> show route snooping detail
__+domainAll__.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
224.0.0.2/32 (1 entry, 1 announced)
  *IGMP    Preference: 0
           Next hop type: MultiRecv
           Next-hop reference count: 4
           State: <Active NoReadvrt Int>
```

```
Age: 2:24
Task: IGMP
Announcement bits (1): 0-KRT
AS path: I

224.0.0.22/32 (1 entry, 1 announced)
  *IGMP Preference: 0
    Next hop type: MultiRecv
    Next-hop reference count: 4
    State: <Active NoReadvrt Int>
    Age: 2:24
    Task: IGMP
    Announcement bits (1): 0-KRT
    AS path: I

__+domainAll__.inet.1: 36 destinations, 36 routes (36 active, 0 holddown, 0 hidden)

224.0.0.0.0.0.0.0.0/24 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4), Next hop index: 1048584
    Next-hop reference count: 4
    State: <Active Int>
    Age: 2:24
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.2.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:13
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.3.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.4.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:17
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.5.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
```



```
State: <Active Int>
Age: 1:58
Task: MC
Announcement bits (1): 0-KRT
AS path: I

225.0.0.6.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:14
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.7.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:12
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.9.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:13
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

225.0.0.10.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

226.0.0.1.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:09
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

226.0.0.2.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
```

```
Age: 8
Task: MC
Announcement bits (1): 0-KRT
AS path: I

226.0.0.4.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:10
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

226.0.0.8.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:12
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

226.0.0.10.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 1:56
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.1.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:10
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.2.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:13
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.3.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:16
```

```
Task: MC
Announcement bits (1): 0-KRT
AS path: I

227.0.0.4.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.5.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 1:57
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.7.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 1:57
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.8.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:10
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

227.0.0.10.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.1.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:09
    Task: MC
```

```
Announcement bits (1): 0-KRT
AS path: I

228.0.0.2.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:18
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.7.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:11
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.8.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:17
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.9.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 8
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

228.0.0.10.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:12
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.3.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:09
    Task: MC
    Announcement bits (1): 0-KRT
```

```
AS path: I

229.0.0.4.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:12
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.5.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 9
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.6.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.7.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.8.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:15
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I

229.0.0.9.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:14
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I
```

```
229.0.0.10.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
  *Multicast Preference: 180
    Next hop type: Multicast (IPv4)
    Next-hop reference count: 113
    State: <Active Int>
    Age: 2:13
    Task: MC
    Announcement bits (1): 0-KRT
    AS path: I
```

show route source-gateway

Syntax	<code>show route source-gateway <i>address</i></code> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	<code>show route source-gateway <i>address</i></code> <brief detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the entries in the routing table that were learned from a particular address. The Source field in the <code>show route detail</code> command output lists the source for each route, if known.
Options	brief detail extensive terse —(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief . <i>address</i> —IP address of the system. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
List of Sample Output	show route source-gateway on page 941 show route source-gateway detail on page 942 show route source-gateway extensive on page 944
Output Fields	For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route source-gateway

```

user@host> show route source-gateway 10.255.70.103
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

```

```
inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:1:3:1/96
    *[BGP/170] 12:12:24, localpref 100, from 10.255.70.103
    AS path: I
    > via so-0/3/0.0, label-switched-path green-r1-r3

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:2:3:1/96
    *[BGP/170] 12:12:24, localpref 0, from 10.255.70.103
    AS path: I
    > via so-0/3/0.0, label-switched-path green-r1-r3

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:1:3:1/96
    *[BGP/170] 12:12:24, localpref 100, from 10.255.70.103
    AS path: I
    > via so-0/3/0.0, label-switched-path green-r1-r3

10.255.70.103:2:3:1/96
    *[BGP/170] 12:12:24, localpref 0, from 10.255.70.103
    AS path: I
    > via so-0/3/0.0, label-switched-path green-r1-r3
```

show route source-gateway detail

```
user@host> show route source-gateway 10.255.70.103 detail
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete
green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

Restart Complete
10.255.70.103:1:3:1/96 (1 entry, 1 announced)
    *BGP    Preference: 170/-101
```



```

Route Distinguisher: 10.255.70.103:1
Next-hop reference count: 7
Source: 10.255.70.103
Protocol next hop: 10.255.70.103
Indirect next hop: 2 no-forward
State: <Secondary Active Int Ext>
Local AS: 69 Peer AS: 69
Age: 12:14:00 Metric2: 1
Task: BGP_69.10.255.70.103+179
Announcement bits (1): 0-green-l2vpn
AS path: I
Communities: target:11111:1 Layer2-info: encaps:VPLS,
control flags:, mtu: 0
Label-base: 800008, range: 8
Localpref: 100
Router ID: 10.255.70.103
Primary Routing Table bgp.l2vpn.0

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:2:3:1/96 (1 entry, 1 announced)
*BGP Preference: 170/-1
Route Distinguisher: 10.255.70.103:2
Next-hop reference count: 7
Source: 10.255.70.103
Protocol next hop: 10.255.70.103
Indirect next hop: 2 no-forward
State: <Secondary Active Int Ext>
Local AS: 69 Peer AS: 69
Age: 12:14:00 Metric2: 1
Task: BGP_69.10.255.70.103+179
Announcement bits (1): 0-red-l2vpn
AS path: I
Communities: target:11111:2 Layer2-info: encaps:VPLS,
control flags:Site-Down, mtu: 0
Label-base: 800016, range: 8
Localpref: 0
Router ID: 10.255.70.103
Primary Routing Table bgp.l2vpn.0

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:1:3:1/96 (1 entry, 0 announced)
*BGP Preference: 170/-101
Route Distinguisher: 10.255.70.103:1
Next-hop reference count: 7
Source: 10.255.70.103
Protocol next hop: 10.255.70.103
Indirect next hop: 2 no-forward
State: <Active Int Ext>
Local AS: 69 Peer AS: 69
Age: 12:14:00 Metric2: 1
Task: BGP_69.10.255.70.103+179
AS path: I
Communities: target:11111:1 Layer2-info: encaps:VPLS, control
flags:, mtu: 0
Label-base: 800008, range: 8
Localpref: 100
Router ID: 10.255.70.103

```

```

Secondary Tables: green.l2vpn.0
10.255.70.103:2:3:1/96 (1 entry, 0 announced)
  *BGP Preference: 170/-1
    Route Distinguisher: 10.255.70.103:2
    Next-hop reference count: 7
    Source: 10.255.70.103
    Protocol next hop: 10.255.70.103
    Indirect next hop: 2 no-forward
    State: <Active Int Ext>
    Local AS: 69 Peer AS: 69
    Age: 12:14:00 Metric2: 1
    Task: BGP_69.10.255.70.103+179
    AS path: I
    Communities: target:11111:2 Layer2-info: encaps:VPLS,
    control flags:Site-Down,
    mtu: 0
    Label-base: 800016, range: 8
    Localpref: 0
    Router ID: 10.255.70.103
    Secondary Tables: red.l2vpn.0

```

show route source-gateway extensive

```

user@host> show route source-gateway 10.255.70.103 extensive
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1___.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
10.255.70.103:1:3:1/96 (1 entry, 1 announced)
  *BGP Preference: 170/-101
    Route Distinguisher: 10.255.70.103:1
    Next-hop reference count: 7
    Source: 10.255.70.103
    Protocol next hop: 10.255.70.103
    Indirect next hop: 2 no-forward
    State: <Secondary Active Int Ext>
    Local AS: 69 Peer AS: 69
    Age: 12:15:24 Metric2: 1
    Task: BGP_69.10.255.70.103+179
    Announcement bits (1): 0-green-l2vpn
    AS path: I
    Communities: target:11111:1 Layer2-info: encaps:VPLS,
    control flags:, mtu: 0
    Label-base: 800008, range: 8
    Localpref: 100
    Router ID: 10.255.70.103

```

Primary Routing Table bgp.l2vpn.0

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:2:3:1/96 (1 entry, 1 announced)

```
*BGP      Preference: 170/-1
          Route Distinguisher: 10.255.70.103:2
          Next-hop reference count: 7
          Source: 10.255.70.103
          Protocol next hop: 10.255.70.103
          Indirect next hop: 2 no-forward
          State: <Secondary Active Int Ext>
          Local AS: 69 Peer AS: 69
          Age: 12:15:24 Metric2: 1
          Task: BGP_69.10.255.70.103+179
          Announcement bits (1): 0-red-l2vpn
          AS path: I
          Communities: target:11111:2 Layer2-info: encaps:VPLS,
          control flags:Site-Down, mtu: 0
          Label-base: 800016, range: 8
          Localpref: 0
          Router ID: 10.255.70.103
          Primary Routing Table bgp.l2vpn.0
```

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:1:3:1/96 (1 entry, 0 announced)

```
*BGP      Preference: 170/-101
          Route Distinguisher: 10.255.70.103:1
          Next-hop reference count: 7
          Source: 10.255.70.103
          Protocol next hop: 10.255.70.103
          Indirect next hop: 2 no-forward
          State: <Active Int Ext>
          Local AS: 69 Peer AS: 69
          Age: 12:15:24 Metric2: 1
          Task: BGP_69.10.255.70.103+179
          AS path: I
          Communities: target:11111:1 Layer2-info: encaps:VPLS,
          control flags:, mtu: 0
          Label-base: 800008, range: 8
          Localpref: 100
          Router ID: 10.255.70.103
          Secondary Tables: green.l2vpn.0
          Indirect next hops: 1
            Protocol next hop: 10.255.70.103 Metric: 2
            Indirect next hop: 2 no-forward
            Indirect path forwarding next hops: 1
          Next hop:      via so-0/3/0.0 weight 0x1
            10.255.70.103/32 Originating RIB: inet.3
              Metric: 2 Node path count: 1
              Forwarding nexthops: 1
                Nexthop: via so-0/3/0.0
```

10.255.70.103:2:3:1/96 (1 entry, 0 announced)

```
*BGP      Preference: 170/-1
          Route Distinguisher: 10.255.70.103:2
          Next-hop reference count: 7
          Source: 10.255.70.103
```

```
Protocol next hop: 10.255.70.103
Indirect next hop: 2 no-forward
State: <Active Int Ext>
Local AS: 69 Peer AS: 69
Age: 12:15:24 Metric2: 1
Task: BGP_69.10.255.70.103+179
AS path: I
Communities: target:11111:2 Layer2-info: encaps:VPLS,
control flags:Site-Down,
mtu: 0
Label-base: 800016, range: 8
Localpref: 0
Router ID: 10.255.70.103
Secondary Tables: red.12vpn.0
Indirect next hops: 1
    Protocol next hop: 10.255.70.103 Metric: 2
    Indirect next hop: 2 no-forward
    Indirect path forwarding next hops: 1
Next hop:      via so-0/3/0.0 weight 0x1
    10.255.70.103/32 Originating RIB: inet.3
    Metric: 2 Node path count: 1
    Forwarding nexthops: 1
    Nexthop: via so-0/3/0.0
```

show route summary

Syntax	show route summary <logical-system (all <i>logical-system-name</i>)> <table <i>routing-table-name</i> >
Syntax (EX Series Switches)	show route summary
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Display summary statistics about the entries in the routing table.</p> <p>CPU utilization might increase while the device learns routes. We recommend that you use the show route summary command after the device learns and enters the routes into the routing table. Depending on the size of your network, this might take several minutes. If you receive a “timeout communicating with routing daemon” error when using the show route summary command, wait several minutes before attempting to use the command again. This is not a critical system error, but you might experience a delay in using the command-line interface (CLI).</p>
Options	<p>none—Display summary statistics about the entries in the routing table.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>table <i>routing-table-name</i>—(Optional) Display summary statistics for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the show route summary table inet command). If you only want to display statistics for a specific routing table, make sure to enter the exact name of that routing table.</p>
Required Privilege Level	view
List of Sample Output	show route summary on page 948 show route summary table on page 949 show route summary table (with Route Limits Configured for the Routing Table) on page 949
Output Fields	Table 141 on page 947 lists the output fields for the show route summary command. Output fields are listed in the approximate order in which they appear.

Table 141: show route summary Output Fields

Field Name	Field Description
Router ID	Address of the local routing device.
<i>routing-table-name</i>	Name of the routing table (for example, inet.0).

Table 141: show route summary Output Fields (*continued*)

Field Name	Field Description
destinations	Number of destinations for which there are routes in the routing table.
routes	Number of routes in the routing table: <ul style="list-style-type: none"> • active—Number of routes that are active. • holddown—Number of routes that are in the hold-down state before being declared inactive. • hidden—Number of routes that are not used because of routing policy.
Limit/Threshold	Displays the configured route limits for the routing table set with the maximum-prefixes and the maximum-paths statements. If you do not configure route limits for the routing table, the show output does not display this information. <ul style="list-style-type: none"> • destinations—The first number represents the maximum number of route prefixes installed in the routing table. The second number represents the number of route prefixes that trigger a warning message. • routes—The first number represents the maximum number of routes. The second number represents the number of routes that trigger a warning message.
Direct	Routes on the directly connected network.
Local	Local routes.
protocol-name	Name of the protocol from which the route was learned. For example, OSPF, RSVP, and Static.

Sample Output

show route summary

```

user@host> show route summary
Autonomous system number: 69
Router ID: 10.255.71.52
Maximum-ECMP: 32
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
      Direct:    6 routes,      5 active
      Local:    4 routes,      4 active
      OSPF:     5 routes,      4 active
      Static:   7 routes,      7 active
      IGMP:     1 routes,      1 active
      PIM:      2 routes,      2 active

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
      RSVP:      2 routes,      2 active

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete
      Direct:    1 routes,      1 active

```

```

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
      MPLS:      3 routes,      3 active
      VPLS:      4 routes,      2 active

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete
      Direct:    2 routes,      2 active
      PIM:       2 routes,      2 active
      MLD:       1 routes,      1 active

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
      BGP:       2 routes,      2 active
      L2VPN:     2 routes,      2 active

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete
      BGP:       2 routes,      2 active
      L2VPN:     1 routes,      1 active

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
      BGP:       4 routes,      4 active

```

show route summary table

```

user@host> show route summary table inet
Router ID: 192.168.0.1

inet.0: 32 destinations, 34 routes (31 active, 0 holddown, 1 hidden)
      Direct:    6 routes,      5 active
      Local:     9 routes,      9 active
      OSPF:      3 routes,      1 active
      Static:    13 routes,     13 active
      IGMP:      1 routes,      1 active
      PIM:       2 routes,      2 active

inet.1: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
      Multicast:  1 routes,      1 active

inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
      Local:     1 routes,      1 active
      PIM:       2 routes,      2 active

inet6.1: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
      Multicast:  1 routes,      1 active

```

show route summary table (with Route Limits Configured for the Routing Table)

```

user@host> show route summary table VPN-A.inet.0
Autonomous system number: 100
Router ID: 10.255.182.142

VPN-A.inet.0: 13 destinations, 14 routes (13 active, 0 holddown, 0 hidden)
Limit/Threshold: 2000/200 destinations 20/12 routes
      Direct:    2 routes,      2 active
      Local:     1 routes,      1 active
      OSPF:      4 routes,      3 active
      BGP:       4 routes,      4 active

```

IGMP:	1 routes,	1 active
PIM:	2 routes,	2 active

show route table

Syntax	<pre>show route table <i>routing-table-name</i> <brief detail extensive terse> <logical-system (all <i>logical-system-name</i>)></pre>
Syntax (EX Series Switches)	<pre>show route table <i>routing-table-name</i> <brief detail extensive terse></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	Display the route entries in a particular routing table.
Options	<p>brief detail extensive terse—(Optional) Display the specified level of output.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>routing-table-name</i>—Display route entries for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the show route table inet command).</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show route summary on page 947
List of Sample Output	<p>show route table bgp.l2.vpn on page 952</p> <p>show route table bgp.l3vpn.0 on page 952</p> <p>show route table bgp.l3vpn.0 detail on page 952</p> <p>show route table bgp.rtarget.0 (When Proxy BGP Route Target Filtering Is Configured) on page 954</p> <p>show route table inet.0 on page 954</p> <p>show route table inet6.0 on page 954</p> <p>show route table inet6.3 on page 955</p> <p>show route table inetflow detail on page 955</p> <p>show route table l2circuit.0 on page 955</p> <p>show route table mpls on page 956</p> <p>show route table mpls extensive on page 956</p> <p>show route table mpls.0 on page 956</p> <p>show route table mpls.0 detail (PTX Series) on page 957</p> <p>show route table mpls.0 extensive (PTX Series) on page 957</p> <p>show route table mpls.0 (RSVP Route—Transit LSP) on page 958</p> <p>show route table vpls_1 detail on page 958</p> <p>show route table vpn-a on page 959</p> <p>show route table vpn-a.mdt.0 on page 959</p> <p>show route table VPN-A detail on page 959</p> <p>show route table VPN-AB.inet.0 on page 960</p>

[show route table VPN_blue.mvpn-inet6.0 on page 960](#)

[show route table vrf1.mvpn.0 extensive on page 961](#)

[show route table inetflow detail on page 961](#)

Output Fields For information about output fields, see the output field tables for the [show route](#) command, the [show route detail](#) command, the [show route extensive](#) command, or the [show route terse](#) command.

Sample Output

show route table bgp.l2.vpn

```
user@host> show route table bgp.l2.vpn
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.24.1:1:4:1/96
    *[BGP/170] 01:08:58, localpref 100, from 192.168.24.1
    AS path: I
    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
```

show route table bgp.l3vpn.0

```
user@host> show route table bgp.l3vpn.0
bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.71.15:100:10.255.71.17/32
    *[BGP/170] 00:03:59, MED 1, localpref 100, from
10.255.71.15
    AS path: I
    > via so-2/1/0.0, Push 100020, Push 100011(top)
10.255.71.15:200:10.255.71.18/32
    *[BGP/170] 00:03:59, MED 1, localpref 100, from
10.255.71.15
    AS path: I
    > via so-2/1/0.0, Push 100021, Push 100011(top)
```

show route table bgp.l3vpn.0 detail

```
user@host> show route table bgp.l3vpn.0 detail
bgp.l3vpn.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)

10.255.245.12:1:4.0.0.0/8 (1 entry, 1 announced)
  *BGP Preference: 170/-101
  Route Distinguisher: 10.255.245.12:1
  Source: 10.255.245.12
  Next hop: 192.168.208.66 via fe-0/0/0.0, selected
  Label operation: Push 182449
  Protocol next hop: 10.255.245.12
  Push 182449
  Indirect next hop: 863a630 297
  State: <Active Int Ext>
  Local AS: 35 Peer AS: 35
  Age: 12:19 Metric2: 1
  Task: BGP_35.10.255.245.12+179
  Announcement bits (1): 0-BGP.0.0.0.0+179
  AS path: 30 10458 14203 2914 3356 I (Atomic) Aggregator: 3356 4.68.0.11

  Communities: 2914:420 target:11111:1 origin:56:78
```

```

VPN Label: 182449
Localpref: 100
Router ID: 10.255.245.12

10.255.245.12:1:4.17.225.0/24 (1 entry, 1 announced)
*BGP Preference: 170/-101
Route Distinguisher: 10.255.245.12:1
Source: 10.255.245.12
Next hop: 192.168.208.66 via fe-0/0/0.0, selected
Label operation: Push 182465
Protocol next hop: 10.255.245.12
Push 182465
Indirect next hop: 863a8f0 305
State: <Active Int Ext>
Local AS: 35 Peer AS: 35
Age: 12:19 Metric2: 1
Task: BGP_35.10.255.245.12+179
Announcement bits (1): 0-BGP.0.0.0.0+179
AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496 I
Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
VPN Label: 182465
Localpref: 100
Router ID: 10.255.245.12

10.255.245.12:1:4.17.226.0/23 (1 entry, 1 announced)
*BGP Preference: 170/-101
Route Distinguisher: 10.255.245.12:1
Source: 10.255.245.12
Next hop: 192.168.208.66 via fe-0/0/0.0, selected
Label operation: Push 182465
Protocol next hop: 10.255.245.12
Push 182465
Indirect next hop: 86bd210 330
State: <Active Int Ext>
Local AS: 35 Peer AS: 35
Age: 12:19 Metric2: 1
Task: BGP_35.10.255.245.12+179
Announcement bits (1): 0-BGP.0.0.0.0+179
AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496
6496 I
Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
VPN Label: 182465
Localpref: 100
Router ID: 10.255.245.12

10.255.245.12:1:4.17.251.0/24 (1 entry, 1 announced)
*BGP Preference: 170/-101
Route Distinguisher: 10.255.245.12:1
Source: 10.255.245.12
Next hop: 192.168.208.66 via fe-0/0/0.0, selected
Label operation: Push 182465
Protocol next hop: 10.255.245.12
Push 182465
Indirect next hop: 86bd210 330
State: <Active Int Ext>
Local AS: 35 Peer AS: 35
Age: 12:19 Metric2: 1
Task: BGP_35.10.255.245.12+179
Announcement bits (1): 0-BGP.0.0.0.0+179
AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496

```

```
6496 I
Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
VPN Label: 182465
Localpref: 100
```

show route table bgp.rtarget.0 (When Proxy BGP Route Target Filtering Is Configured)

```
user@host> show route table bgp.rtarget.0
bgp.rtarget.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100:100:100/96
* [RTarget/5] 00:03:14
  Type Proxy
    for 10.255.165.103
    for 10.255.166.124
  Local
```

show route table inet.0

```
user@host> show route table inet.0
inet.0: 12 destinations, 12 routes (11 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0 * [Static/5] 00:51:57
          > to 111.222.5.254 via fxp0.0
1.0.0.1/32 * [Direct/0] 00:51:58
          > via at-5/3/0.0
1.0.0.2/32 * [Local/0] 00:51:58
          Local
12.12.12.21/32 * [Local/0] 00:51:57
          Reject
13.13.13.13/32 * [Direct/0] 00:51:58
          > via t3-5/2/1.0
13.13.13.14/32 * [Local/0] 00:51:58
          Local
13.13.13.21/32 * [Local/0] 00:51:58
          Local
13.13.13.22/32 * [Direct/0] 00:33:59
          > via t3-5/2/0.0
127.0.0.1/32 [Direct/0] 00:51:58
          > via lo0.0
111.222.5.0/24 * [Direct/0] 00:51:58
          > via fxp0.0
111.222.5.81/32 * [Local/0] 00:51:58
          Local
```

show route table inet6.0

```
user@host> show route table inet6.0
inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Route, * = Both

fec0:0:0:3::/64 * [Direct/0] 00:01:34
> via fe-0/1/0.0

fec0:0:0:3::/128 * [Local/0] 00:01:34
> Local

fec0:0:0:4::/64 * [Static/5] 00:01:34
> to fec0:0:0:3::ffff via fe-0/1/0.0
```

show route table inet6.3

```

user@router> show route table inet6.3
inet6.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::10.255.245.195/128
    *[LDP/9] 00:00:22, metric 1
    > via so-1/0/0.0
::10.255.245.196/128
    *[LDP/9] 00:00:08, metric 1
    > via so-1/0/0.0, Push 100008

```

show route table inetflow detail

```

user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
    *BGP Preference: 170/-101
        Next-hop reference count: 2
        State: <Active Ext>
        Local AS: 65002 Peer AS: 65000
        Age: 4
        Task: BGP_65000.10.12.99.5+3792
        Announcement bits (1): 0-Flow
        AS path: 65000 I
        Communities: traffic-rate:0:0
        Validation state: Accept, Originator: 10.12.99.5
        Via: 10.12.44.0/24, Active
        Localpref: 100
        Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
    *Flow Preference: 5
        Next-hop reference count: 2
        State: <Active>
        Local AS: 65002
        Age: 6:30
        Task: RT Flow
        Announcement bits (2): 0-Flow 1-BGP.0.0.0.0+179
        AS path: I
        Communities: 1:1

```

show route table l2circuit.0

```

user@host> show route table l2circuit.0
l2circuit.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.195:NoCtrlWord:1:1:Local/96
    *[L2CKT/7] 00:50:47
    > via so-0/1/2.0, Push 100049
    > via so-0/1/3.0, Push 100049
10.1.1.195:NoCtrlWord:1:1:Remote/96
    *[LDP/9] 00:50:14
    Discard
10.1.1.195:CtrlWord:1:2:Local/96
    *[L2CKT/7] 00:50:47
    > via so-0/1/2.0, Push 100049
    > via so-0/1/3.0, Push 100049
10.1.1.195:CtrlWord:1:2:Remote/96

```

```
*[LDP/9] 00:50:14
Discard
```

show route table mpls

```
user@host> show route table mpls
mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                *[MPLS/0] 00:13:55, metric 1
                  Receive
1                *[MPLS/0] 00:13:55, metric 1
                  Receive
2                *[MPLS/0] 00:13:55, metric 1
                  Receive
1024             *[VPN/0] 00:04:18
                  to table red.inet.0, Pop
```

show route table mpls extensive

```
user@host> show route table mpls extensive
100000 (1 entry, 1 announced)
TSI:
KRT in-kernel 100000 /36 -> {so-1/0/0.0}
    *LDP      Preference: 9
              Next hop: via so-1/0/0.0, selected
              Pop
              State: <Active Int>
              Age: 29:50      Metric: 1
              Task: LDP
              Announcement bits (1): 0-KRT
              AS path: I
              Prefixes bound to route: 10.0.0.194/32
```

show route table mpls.0

```
user@host> show route table mpls.0
mpls.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                *[MPLS/0] 00:45:09, metric 1
                  Receive
1                *[MPLS/0] 00:45:09, metric 1
                  Receive
2                *[MPLS/0] 00:45:09, metric 1
                  Receive
100000           *[L2VPN/7] 00:43:04
                  > via so-0/1/0.1, Pop
100001           *[L2VPN/7] 00:43:03
                  > via so-0/1/0.2, Pop      Offset: 4
100002           *[LDP/9] 00:43:22, metric 1
                  via so-0/1/2.0, Pop
                  > via so-0/1/3.0, Pop
100002(S=0)      *[LDP/9] 00:43:22, metric 1
                  via so-0/1/2.0, Pop
                  > via so-0/1/3.0, Pop
100003           *[LDP/9] 00:43:22, metric 1
                  > via so-0/1/2.0, Swap 100002
                  via so-0/1/3.0, Swap 100002
100004           *[LDP/9] 00:43:16, metric 1
                  via so-0/1/2.0, Swap 100049
                  > via so-0/1/3.0, Swap 100049
```

```

so-0/1/0.1      *[L2VPN/7] 00:43:04
                 > via so-0/1/2.0, Push 100001, Push 100049(top)
                 via so-0/1/3.0, Push 100001, Push 100049(top)
so-0/1/0.2      *[L2VPN/7] 00:43:03
                 via so-0/1/2.0, Push 100000, Push 100049(top) Offset: -4
                 > via so-0/1/3.0, Push 100000, Push 100049(top) Offset: -4

```

show route table mpls.0 detail (PTX Series)

```

user@host> show route table mpls.0 detail
ge-0/0/2.600 (1 entry, 1 announced)
  *L2VPN Preference: 7
    Next hop type: Indirect
    Address: 0x9438f34
    Next-hop reference count: 2
    Next hop type: Router, Next hop index: 567
    Next hop: 3.0.0.1 via ge-0/0/1.0, selected
    Label operation: Push 299808
    Label TTL action: prop-ttl
    Load balance label: Label 299808:None;
    Session Id: 0x1
    Protocol next hop: 10.255.255.1
    Label operation: Push 299872 Offset: 252
    Label TTL action: no-prop-ttl
    Load balance label: Label 299872:Flow label PUSH;
    Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
    Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
    State: <Active Int>
    Age: 21 Metric2: 1
    Validation State: unverified
    Task: Common L2 VC
    Announcement bits (2): 0-KRT 2-Common L2 VC
    AS path: I

```

show route table mpls.0 extensive (PTX Series)

```

user@host> show route table mpls.0 extensive
ge-0/0/2.600 (1 entry, 1 announced)
TSI:
KRT in-kernel ge-0/0/2.600.0 /32 -> {composite(570)}
  *L2VPN Preference: 7
    Next hop type: Indirect
    Address: 0x9438f34
    Next-hop reference count: 2
    Next hop type: Router, Next hop index: 567
    Next hop: 3.0.0.1 via ge-0/0/1.0, selected
    Label operation: Push 299808
    Label TTL action: prop-ttl
    Load balance label: Label 299808:None;
    Session Id: 0x1
    Protocol next hop: 10.255.255.1
    Label operation: Push 299872 Offset: 252
    Label TTL action: no-prop-ttl
    Load balance label: Label 299872:Flow label PUSH;
    Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
    Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
    State: <Active Int>
    Age: 47 Metric2: 1
    Validation State: unverified
    Task: Common L2 VC
    Announcement bits (2): 0-KRT 2-Common L2 VC

```

```

AS path: I
Composite next hops: 1
  Protocol next hop: 10.255.255.1 Metric: 1
  Label operation: Push 299872 Offset: 252
  Label TTL action: no-prop-ttl
  Load balance label: Label 299872:Flow label PUSH;
  Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
  Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
  Indirect path forwarding next hops: 1
    Next hop type: Router
    Next hop: 3.0.0.1 via ge-0/0/1.0
    Session Id: 0x1
  10.255.255.1/32 Originating RIB: inet.3
  Metric: 1 Node path count: 1
  Forwarding nexthops: 1
  Nexthop: 3.0.0.1 via ge-0/0/1.0

```

show route table mpls.0 (RSVP Route—Transit LSP)

```
user@host> show route table mpls.0
```

```

mpls.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

```

```

0          *[MPLS/0] 00:37:31, metric 1
            Receive
1          *[MPLS/0] 00:37:31, metric 1
            Receive
2          *[MPLS/0] 00:37:31, metric 1
            Receive
13         *[MPLS/0] 00:37:31, metric 1
            Receive
300352     *[RSVP/7/1] 00:08:00, metric 1
            > to 8.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300352(S=0) *[RSVP/7/1] 00:08:00, metric 1
            > to 8.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300384     *[RSVP/7/2] 00:05:20, metric 1
            > to 8.64.1.106 via ge-1/0/0.0, Pop
300384(S=0) *[RSVP/7/2] 00:05:20, metric 1
            > to 8.64.1.106 via ge-1/0/0.0, Pop

```

show route table vpls_1 detail

```
user@host> show route table vpls_1 detail
```

```

vpls_1.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

```

```

1.1.1.11:1000:1:1/96 (1 entry, 1 announced)
*L2VPN Preference: 170/-1
Receive table: vpls_1.l2vpn.0
Next-hop reference count: 2
State: <Active Int Ext>
Age: 4:29:47 Metric2: 1
Task: vpls_1-l2vpn
Announcement bits (1): 1-BGP.0.0.0.0+179
AS path: I
Communities: Layer2-info: encaps:VPLS, control flags:Site-Down
Label-base: 800000, range: 8, status-vector: 0xFF

```


show route table vpn-a

```

user@host> show route table vpn-a
vpn-a.12vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both
192.168.16.1:1:1:1/96
    *[VPN/7] 05:48:27
    Discard
192.168.24.1:1:2:1/96
    *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
    AS path: I
    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
192.168.24.1:1:3:1/96
    *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
    AS path: I
    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am

```

show route table vpn-a.mdt.0

```

user@host> show route table vpn-a.mdt.0
vpn-a.mdt.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:1:0:10.255.14.216:232.1.1.1/144
    *[MVPN/70] 01:23:05, metric2 1
    Indirect
1:1:1:10.255.14.218:232.1.1.1/144
    *[BGP/170] 00:57:49, localpref 100, from 10.255.14.218
    AS path: I
    > via so-0/0/0.0, label-switched-path r0e-to-r1
1:1:2:10.255.14.217:232.1.1.1/144
    *[BGP/170] 00:57:49, localpref 100, from 10.255.14.217
    AS path: I
    > via so-0/0/1.0, label-switched-path r0-to-r2

```

show route table VPN-A detail

```

user@host> show route table VPN-A detail
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
10.255.179.9/32 (1 entry, 1 announced)
    *BGP
    Preference: 170/-101
    Route Distinguisher: 10.255.179.13:200
    Next hop type: Indirect
    Next-hop reference count: 5
    Source: 10.255.179.13
    Next hop type: Router, Next hop index: 732
    Next hop: 10.39.1.14 via fe-0/3/0.0, selected
    Label operation: Push 299824, Push 299824(top)
    Protocol next hop: 10.255.179.13
    Push 299824
    Indirect next hop: 8f275a0 1048574
    State: (Secondary Active Int Ext)
    Local AS: 1 Peer AS: 1
    Age: 3:41:06 Metric: 1 Metric2: 1
    Task: BGP_1.10.255.179.13+64309
    Announcement bits (2): 0-KRT 1-BGP RT Background
    AS path: I
    Communities: target:1:200 rte-type:0.0.0.0:1:0
    Import Accepted
    VPN Label: 299824 TTL Action: vrf-ttl-propagate

```

```

Localpref: 100
Router ID: 10.255.179.13
Primary Routing Table bgp.13vpn.0

```

show route table VPN-AB.inet.0

```

user@host> show route table VPN-AB.inet.0
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.0/30      *[OSPF/10] 00:07:24, metric 1
                  > via so-7/3/1.0
10.39.1.4/30      *[Direct/0] 00:08:42
                  > via so-5/1/0.0
10.39.1.6/32      *[Local/0] 00:08:46
                  Local
10.255.71.16/32   *[Static/5] 00:07:24
                  > via so-2/0/0.0
10.255.71.17/32   *[BGP/170] 00:07:24, MED 1, localpref 100, from
10.255.71.15
                  AS path: I
                  > via so-2/1/0.0, Push 100020, Push 100011(top)
10.255.71.18/32   *[BGP/170] 00:07:24, MED 1, localpref 100, from
10.255.71.15
                  AS path: I
                  > via so-2/1/0.0, Push 100021, Push 100011(top)
10.255.245.245/32 *[BGP/170] 00:08:35, localpref 100
                  AS path: 2 I
                  > to 10.39.1.5 via so-5/1/0.0
10.255.245.246/32 *[OSPF/10] 00:07:24, metric 1
                  > via so-7/3/1.0

```

show route table VPN_blue.mvpn-inet6.0

```

user@host> show route table VPN_blue.mvpn-inet6.0
vpn_blue.mvpn-inet6.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:10.255.2.202:65535:10.255.2.202/432
                  *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
                  AS path: I
                  > via so-0/1/3.0
1:10.255.2.203:65535:10.255.2.203/432
                  *[BGP/170] 00:02:37, localpref 100, from 10.255.2.203
                  AS path: I
                  > via so-0/1/0.0
1:10.255.2.204:65535:10.255.2.204/432
                  *[MVPN/70] 00:57:23, metric2 1
                  Indirect
5:10.255.2.202:65535:128::192.168.90.2:128:ffff::1/432
                  *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
                  AS path: I
                  > via so-0/1/3.0
6:10.255.2.203:65535:65000:128::10.12.53.12:128:ffff::1/432
                  *[PIM/105] 00:02:37
                  Multicast (IPv6)
7:10.255.2.202:65535:65000:128::192.168.90.2:128:ffff::1/432
                  *[MVPN/70] 00:02:37, metric2 1
                  Indirect

```

show route table vrf1.mvpn.0 extensive

```

user@host> show route table vrf1.mvpn.0 extensive
1:10.255.50.77:1:10.255.50.77/240 (1 entry, 1 announced)
    *MVPN    Preference: 70
              PMSI: Flags 0x0: Label 0: RSVP-TE:
Session_13[10.255.50.77:0:25624:10.255.50.77]
    Next hop type: Indirect
    Address: 0xbb2c944
    Next-hop reference count: 360
    Protocol next hop: 10.255.50.77
    Indirect next hop: 0x0 - INH Session ID: 0x0
    State: <Active Int Ext>
    Age: 53:03      Metric2: 1
    Validation State: unverified
    Task: mvpn global task
    Announcement bits (3): 0-PIM.vrf1 1-mvpn global task 2-rt-export

    AS path: I

```

show route table inetflow detail

```

user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
    *BGP      Preference: 170/-101
              Next-hop reference count: 2
              State: <Active Ext>
              Local AS: 65002 Peer AS: 65000
              Age: 4
              Task: BGP_65000.10.12.99.5+3792
              Announcement bits (1): 0-Flow
              AS path: 65000 I
              Communities: traffic-rate:0:0
              Validation state: Accept, Originator: 10.12.99.5
              Via: 10.12.44.0/24, Active
              Localpref: 100
              Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
    *Flow     Preference: 5
              Next-hop reference count: 2
              State: <Active>
              Local AS: 65002
              Age: 6:30
              Task: RT Flow
              Announcement bits (2): 0-Flow 1-BGP.0.0.0.0+179
              AS path: I
              Communities: 1:1

user@PE1> show route table green.l2vpn.0 (VPLS Multihoming with FEC 129)
green.l2vpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.2:100:1.1.1.2/96 AD
    *[VPLS/170] 1d 03:11:03, metric2 1
    Indirect
1.1.1.4:100:1.1.1.4/96 AD
    *[BGP/170] 1d 03:11:02, localpref 100, from 1.1.1.4
    AS path: I, validation-state: unverified
    >    via ge-1/2/1.5

```

```

1.1.1.2:100:1:0/96 MH
    *[VPLS/170] 1d 03:11:03, metric2 1
    Indirect
1.1.1.4:100:1:0/96 MH
    *[BGP/170] 1d 03:11:02, localpref 100, from 1.1.1.4
    AS path: I, validation-state: unverified
    > via ge-1/2/1.5
1.1.1.4:NoCtrlWord:5:100:100:1.1.1.2:1.1.1.4/176
    *[VPLS/7] 1d 03:11:02, metric2 1
    > via ge-1/2/1.5
1.1.1.4:NoCtrlWord:5:100:100:1.1.1.4:1.1.1.2/176
    *[LDP/9] 1d 03:11:02
    Discard

user@host> show route table red extensive
red.inet.0: 364481 destinations, 714087 routes (364480 active, 48448 holddown, 1
hidden)
22.0.0.0/32 (3 entries, 1 announced)
    State: <OnList CalcForwarding>
TSI:
KRT in-kernel 22.0.0.0/32 -> {composite(1048575)} Page 0 idx 1 Type 1 val 0x934342c

    Nexthop: Self
    AS path: [2] I
    Communities: target:2:1
Path 22.0.0.0 from 2.3.0.0 Vector len 4. Val: 1
    @BGP Preference: 170/-1
    Route Distinguisher: 2:1
    Next hop type: Indirect
    Address: 0x258059e4
    Next-hop reference count: 2
    Source: 2.2.0.0
    Next hop type: Router
    Next hop: 10.1.1.1 via ge-1/1/9.0, selected
    Label operation: Push 707633
    Label TTL action: prop-ttl
    Session Id: 0x17d8
    Protocol next hop: 2.2.0.0
    Push 16
    Composite next hop: 0x25805988 - INH Session ID: 0x193c
    Indirect next hop: 0x23eea900 - INH Session ID: 0x193c
    State: <Secondary Active Int Ext ProtectionPath ProtectionCand>
    Local AS: 2 Peer AS: 2
    Age: 23 Metric2: 35
    Validation State: unverified
    Task: BGP_2.2.2.0.0+34549
    AS path: I
    Communities: target:2:1
    Import Accepted
    VPN Label: 16
    Localpref: 0
    Router ID: 2.2.0.0
    Primary Routing Table bgp.13vpn.0
    Composite next hops: 1
        Protocol next hop: 2.2.0.0 Metric: 35
        Push 16
        Composite next hop: 0x25805988 - INH Session ID: 0x193c
        Indirect next hop: 0x23eea900 - INH Session ID: 0x193c
        Indirect path forwarding next hops: 1
            Next hop type: Router
            Next hop: 10.1.1.1 via ge-1/1/9.0

```

```

                Session Id: 0x17d8
                2.2.0.0/32 Originating RIB: inet.3
                Metric: 35                      Node path count: 1
                Forwarding nexthops: 1
                Nexthop: 10.1.1.1 via ge-1/1/9.0
BGP Preference: 170/-1
Route Distinguisher: 2:1
Next hop type: Indirect
Address: 0x9347028
Next-hop reference count: 3
Source: 2.3.0.0
Next hop type: Router, Next hop index: 702
Next hop: 10.1.4.2 via ge-1/0/0.0, selected
Label operation: Push 634278
Label TTL action: prop-ttl
Session Id: 0x17d9
Protocol next hop: 2.3.0.0
Push 16
Composite next hop: 0x93463a0 1048575 INH Session ID: 0x17da
Indirect next hop: 0x91e8800 1048574 INH Session ID: 0x17da
State: <Secondary NotBest Int Ext ProtectionPath ProtectionCand>

Inactive reason: Not Best in its group - IGP metric
Local AS:      2 Peer AS:      2
Age: 3:34      Metric2: 70
Validation State: unverified
Task: BGP_2.2.3.0.0+32805
Announcement bits (2): 0-KRT 1-BGP_RT_Background
AS path: I
Communities: target:2:1
Import Accepted
VPN Label: 16
Localpref: 0
Router ID: 2.3.0.0
Primary Routing Table bgp.13vpn.0
Composite next hops: 1
    Protocol next hop: 2.3.0.0 Metric: 70
    Push 16
    Composite next hop: 0x93463a0 1048575 INH Session ID:
0x17da
    Indirect next hop: 0x91e8800 1048574 INH Session ID:
0x17da
    Indirect path forwarding next hops: 1
        Next hop type: Router
        Next hop: 10.1.4.2 via ge-1/0/0.0
        Session Id: 0x17d9
        2.3.0.0/32 Originating RIB: inet.3
        Metric: 70                      Node path count: 1
        Forwarding nexthops: 1
        Nexthop: 10.1.4.2 via ge-1/0/0.0
#Multipath Preference: 255
Next hop type: Indirect
Address: 0x24afca30
Next-hop reference count: 1
Next hop type: Router
Next hop: 10.1.1.1 via ge-1/1/9.0, selected
Label operation: Push 707633
Label TTL action: prop-ttl
Session Id: 0x17d8
Next hop type: Router, Next hop index: 702
Next hop: 10.1.4.2 via ge-1/0/0.0

```

```

Label operation: Push 634278
Label TTL action: prop-ttl
Session Id: 0x17d9
Protocol next hop: 2.2.0.0
Push 16
Composite next hop: 0x25805988 - INH Session ID: 0x193c
Indirect next hop: 0x23eea900 - INH Session ID: 0x193c Weight 0x1

Protocol next hop: 2.3.0.0
Push 16
Composite next hop: 0x93463a0 1048575 INH Session ID: 0x17da
Indirect next hop: 0x91e8800 1048574 INH Session ID: 0x17da Weight

0x4000
State: <ForwardingOnly Int Ext>
Inactive reason: Forwarding use only
Age: 23          Metric2: 35
Validation State: unverified
Task: RT
AS path: I
Communities: target:2:1

```

show route terse


Syntax	show route terse <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	show route terse
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display a high-level summary of the routes in the routing table.
<div>  <p>NOTE: For BGP routes, the show route terse command displays the local preference attribute and MED instead of the metric1 and metric2 values. This is mostly due to historical reasons.</p> <p>To display the metric1 and metric2 value of a BGP route, use the show route extensive command.</p> </div>	
Options	<p>none—Display a high-level summary of the routes in the routing table.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show route terse on page 967
Output Fields	Table 142 on page 965 describes the output fields for the show route terse command. Output fields are listed in the approximate order in which they appear.

Table 142: show route terse Output Fields

Field Name	Field Description
<i>routing-table-name</i>	Name of the routing table (for example, inet.0).
<i>number destinations</i>	Number of destinations for which there are routes in the routing table.
<i>number routes</i>	Number of routes in the routing table and total number of routes in the following states: <ul style="list-style-type: none"> active (routes that are active) holddown (routes that are in the pending state before being declared inactive) hidden (routes that are not used because of a routing policy)

Table 142: show route terse Output Fields (*continued*)

Field Name	Field Description
<i>route key</i>	Key for the state of the route: <ul style="list-style-type: none"> • +—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • -—A hyphen indicates the last active route. • *—An asterisk indicates that the route is both the active and the last active route. An asterisk before a to line indicates the best subpath to the route.
A	Active route. An asterisk (*) indicates this is the active route.
V	Validation status of the route: <ul style="list-style-type: none"> • ?—Not evaluated. Indicates that the route was not learned through BGP. • I—Invalid. Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • N—Unknown. Indicates that the prefix is not among the prefixes or prefix ranges in the database. • V—Valid. Indicates that the prefix and autonomous system pair are found in the database.
Destination	Destination of the route.
P	Protocol through which the route was learned: <ul style="list-style-type: none"> • A—Aggregate • B—BGP • C—CCC • D—Direct • G—GMPLS • I—IS-IS • L—L2CKT, L2VPN, LDP, Local • K—Kernel • M—MPLS, MSDP • O—OSPF • P—PIM • R—RIP, RIPng • S—Static • T—Tunnel
Prf	Preference value of the route. In every routing metric except for the BGP LocalPref attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the LocalPref value in the Preference2 field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the LocalPref value for Route 2 is 155, the Preference2 value is -156. Route 2 is preferred because it has a higher LocalPref value and a lower Preference2 value.
Metric 1	First metric value in the route. For routes learned from BGP, this is the MED metric.
Metric 2	Second metric value in the route. For routes learned from BGP, this is the IGP metric.

Table 142: show route terse Output Fields (*continued*)

Field Name	Field Description
Next hop	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.
AS path	<p>AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated:</p> <ul style="list-style-type: none"> I—IGP. E—EGP. ?—Incomplete; typically, the AS path was aggregated.

Sample Output

show route terse

```

user@host> show route terse
inet.0: 10 destinations, 12 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A V Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* ? 1.0.1.1/32        0 10      1           >10.0.0.2      I
?                               B 170      100           >10.0.0.2      I
unverified
* ? 1.1.1.1/32        D 0       110         >10.0.0.2      200 I
* V 2.2.0.2/32        B 170      110         >10.0.0.2
valid
* ? 10.0.0.0/30       D 0       100         >1t-1/2/0.1    I
?                               B 170      100           >10.0.0.2
unverified
* ? 10.0.0.1/32       L 0       100         Local          I
* ? 10.0.0.4/30       B 170      100           >10.0.0.2      I
unverified
* ? 10.0.0.8/30       B 170      100           >10.0.0.2      I
unverified
* I 172.16.1.1/32     B 170      90           >10.0.0.2      200 I
invalid
* N 192.168.2.3/32    B 170      100           >10.0.0.2      200 I
unknown
* ? 224.0.0.5/32      O 10      1           MultiRecv

```


CHAPTER 17

RIP Operational Mode Commands

- `clear rip general-statistics`
- `clear rip statistics`
- `show rip general-statistics`
- `show rip neighbor`
- `show rip statistics`

clear rip general-statistics

Syntax	clear rip general-statistics <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and QFX Series)	clear rip general-statistics
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Clear RIP general statistics.
Options	none —Clear RIP general statistics. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show rip general-statistics on page 972
List of Sample Output	clear rip general-statistics on page 970
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear rip general-statistics

```
user@host> clear rip general-statistics
```

clear rip statistics

Syntax	clear rip statistics <instance (all <i>instance-name</i>)> <logical-system (all <i>logical-system-name</i>)> <neighbor> <peer (all <i>address</i>)>
Syntax (EX Series Switches and QFX Series)	clear rip statistics <instance (all <i>instance-name</i>)> <neighbor>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Clear RIP statistics.
Options	<p>none—Reset RIP counters for all neighbors for all routing instances.</p> <p>instance (all <i>instance-name</i>)—(Optional) Clear RIP statistics for all instances or for the specified routing instance only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>neighbor—(Optional) Clear RIP statistics for the specified neighbor only.</p> <p>peer (all <i>address</i>)—(Optional) Clear RIP statistics for a single peer or all peers.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show rip statistics on page 976
List of Sample Output	clear rip statistics on page 971
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear rip statistics

```
user@host> clear rip statistics
```

show rip general-statistics

Syntax	show rip general-statistics <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches and QFX Series)	show rip general-statistics
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display brief RIP statistics.
Options	<p>none—Display brief RIP statistics.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> clear rip general-statistics on page 970
List of Sample Output	show rip general-statistics on page 972
Output Fields	Table 143 on page 972 lists the output fields for the show rip general-statistics command. Output fields are listed in the approximate order in which they appear.

Table 143: show rip general-statistics Output Fields

Field Name	Field Description
bad msgs	Number of invalid messages received.
no rcv intf	Number of packets received with no matching interface.
curr memory	Amount of memory currently used by RIP.
max memory	Most memory used by RIP.

Sample Output

show rip general-statistics

```

user@host> show rip general-statistics
RIPv2 I/O info:
  bad msgs      :      0
  no rcv intf   :      0
  curr memory   :      0
  max memory    :      0

```


show rip neighbor

Syntax	show rip neighbor <instance (all <i>instance-name</i>)> <logical-system (all <i>logical-system-name</i>)> < <i>name</i> >
Syntax (EX Series Switches and QFX Series)	show rip neighbor <instance (all <i>instance-name</i>)> < <i>name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display information about RIP neighbors.
Options	<p>none—Display information about all RIP neighbors for all instances.</p> <p>instance (all <i>instance-name</i>)—(Optional) Display RIP neighbor information for all instances or for only the specified routing instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>name</i>—(Optional) Display detailed information about only the specified RIP neighbor.</p>
Required Privilege Level	view
List of Sample Output	show rip neighbor on page 975 show rip neighbor (With Demand Circuits Configured) on page 975
Output Fields	Table 144 on page 974 lists the output fields for the show rip neighbor command. Output fields are listed in the approximate order in which they appear.

Table 144: show rip neighbor Output Fields

Field Name	Field Description
Neighbor	<p>Name of the RIP neighbor.</p> <p>NOTE: Beginning with Junos OS Release 11.1, when you configure demand circuits, the output displays a demand circuit (DC) flag next to neighbor interfaces configured for demand circuits.</p> <p>If you configure demand circuits at the [edit protocols rip group <i>group-name</i> neighbor <i>neighbor-name</i>] hierarchy level, the output shows only the neighboring interface that you specifically configured as a demand circuit. If you configure demand circuits at the [edit protocols rip group <i>group-name</i>] hierarchy level, all of the interfaces in the group are configured as demand circuits. Therefore, the output shows all of the interfaces in that group as demand circuits.</p>

Table 144: show rip neighbor Output Fields (*continued*)

Field Name	Field Description
State	State of the connection: Up or Dn (Down).
Source Address	Address of the port on the local router.
Destination Address	Address of the port on the remote router.
Send Mode	Send options: broadcast , multicast , none , or version 1 .
Receive Mode	Type of packets to accept: both , none , version 1 , or version 2 .
In Met	Metric added to incoming routes when advertising into RIP routes that were learned from other protocols.

Sample Output

show rip neighbor

```

user@host> show rip neighbor
Neighbor      Local  Source  Destination  Send  Receive  In
-----      -
ge-2/3/0.0    Up    192.168.9.105  192.168.9.107  bcast  both      1
at-5/1/1.42    Dn    (null)      (null)        mcast  v2 only    3
at-5/1/0.42    Dn    (null)      (null)        mcast  both       3
at-5/1/0.0     Up    20.0.0.1     224.0.0.9     mcast  both       3
so-0/0/0.0     Up    192.168.9.97  224.0.0.9     mcast  both       3

```

show rip neighbor (With Demand Circuits Configured)

```

user@host> show rip neighbor
Neighbor      Local  Source  Destination  Send  Receive  In
-----      -
so-0/1/0.0(DC)  Up    10.10.10.2  224.0.0.9     mcast  both       1
so-0/2/0.0(DC)  Up    13.13.13.2  224.0.0.9     mcast  both       1

```

show rip statistics

Syntax	<code>show rip statistics</code> <code><instance (all <i>instance-name</i>)></code> <code><logical-system (all <i>logical-system-name</i>)></code> <code><name></code> <code><peer (all <i>address</i>)></code>
Syntax (EX Series Switches and QFX Series)	<code>show rip statistics</code> <code><instance (all <i>instance-name</i>)></code> <code><name></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 12.1 for the QFX Series.
Description	Display RIP statistics about messages sent and received on an interface, as well as information received from advertisements from other routing devices.
Options	none —Display RIP statistics for all routing instances. instance (all <i>instance-name</i>) —(Optional) Display RIP statistics for all instances or for only the specified routing instance. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. name —(Optional) Display detailed information about only the specified RIP neighbor. peer (all <i>address</i>) —(Optional) Display RIP statistics for a single peer or all peers.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• clear rip statistics on page 971
List of Sample Output	show rip statistics on page 977
Output Fields	Table 145 on page 977 lists the output fields for the show rip statistics command. Output fields are listed in the approximate order in which they appear.

Table 145: show rip statistics Output Fields

Field Name	Field Description
RIP info	<p>Information about RIP on the specified interface:</p> <ul style="list-style-type: none"> • port—UDP port number used for RIP. • update interval—Interval between routing table updates, in seconds. • holddown—Hold-down interval, in seconds. • timeout—Timeout interval, in seconds. • restart in progress—Graceful restart status. Displayed when RIP is or has been in the process of graceful restart. • restart time—Estimated time for the graceful restart to finish, in seconds. • restart will complete in—Remaining time for the graceful restart to finish, in seconds. • rts learned—Number of routes learned through RIP. • rts held down—Number of routes held down by RIP. • rqsts dropped—Number of received request packets that were dropped. • resps dropped—Number of received response packets that were dropped.
logical-interface	<p>Name of the logical interface and its statistics:</p> <ul style="list-style-type: none"> • routes learned—Number of routes learned on the logical interface. • routes advertised—Number of routes advertised by the logical interface.
Counter	<p>List of counter types:</p> <ul style="list-style-type: none"> • Updates Sent—Number of update messages sent. • Triggered Updates Sent—Number of triggered update messages sent. • Responses Sent—Number of response messages sent. • Bad Messages—Number of invalid messages received. • RIPv1 Updates Received—Number of RIPv1 update messages received. • RIPv1 Bad Route Entries—Number of RIPv1 invalid route entry messages received. • RIPv1 Updates Ignored—Number of RIPv1 update messages ignored. • RIPv2 Updates Received—Number of RIPv2 update messages received. • RIPv2 Bad Route Entries—Number of RIPv2 invalid route entry messages received. • RIPv2 Updates Ignored—Number of RIPv2 update messages ignored. • Authentication Failures—Number of received update messages that failed authentication. • RIP Requests Received—Number of RIP request messages received. • RIP Requests Ignored—Number of RIP request messages ignored.
Total	Total number of packets for the selected counter.
Last 5 min	Number of packets for the selected counter in the most recent 5-minute period.
Last minute	Number of packets for the selected counter in the most recent 1-minute period.

Sample Output

show rip statistics

```
user@host> show rip statistics so-0/0/0.0
```

RIP info: port 520; update interval: 30s; holddown 180s; timeout 120s
restart in progress: restart time 60s; restart will complete in 55s
 rts learned rts held down rqsts dropped resps dropped
 0 0 0 0

so-0/0/0.0: 0 routes learned; 501 routes advertised

Counter	Total	Last 5 min	Last minute
-----	-----	-----	-----
Updates Sent	0	0	0
Triggered Updates Sent	0	0	0
Responses Sent	0	0	0
Bad Messages	0	0	0
RIPv1 Updates Received	0	0	0
RIPv1 Bad Route Entries	0	0	0
RIPv1 Updates Ignored	0	0	0
RIPv2 Updates Received	0	0	0
RIPv2 Bad Route Entries	0	0	0
RIPv2 Updates Ignored	0	0	0
Authentication Failures	0	0	0
RIP Requests Received	0	0	0
RIP Requests Ignored	0	0	0

CHAPTER 18

RIPng Operational Mode Commands

- clear ripng general-statistics
- clear ripng statistics
- show ripng general-statistics
- show ripng neighbor
- show ripng statistics

clear ripng general-statistics

Syntax	clear ripng general-statistics <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switches)	clear ripng general-statistics
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Clear RIP next generation (RIPng) general statistics.
Options	none —Clear RIPng general statistics. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show ripng general-statistics on page 982
List of Sample Output	clear ripng general-statistics on page 980
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ripng general-statistics

```
user@host> clear ripng general-statistics
```

clear ripng statistics

Syntax	clear ripng statistics <instance name> <logical-system (all logical-system-name)>
Syntax (EX Series Switch)	clear ripng statistics <instance name>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Clear RIP next-generation (RIPng) statistics.
Options	<p>none—Reset RIPng counters for all neighbors for all routing instances.</p> <p>instance—(Optional) Reset RIPng counters for the specified instance.</p> <p>logical-system (all logical-system-name)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>name—(Optional) Reset RIPng counters for the specified neighbor.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show ripng statistics on page 985
List of Sample Output	clear ripng statistics on page 981
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ripng statistics

```
user@host> clear ripng statistics
```

show ripng general-statistics

Syntax	show ripng general-statistics <logical-system (all <i>logical-system-name</i>)>
Syntax (EX Series Switch)	show ripng general-statistics
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display general RIP next-generation (RIPng) statistics.
Options	none —Display general RIPng statistics. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> clear ripng general-statistics on page 980
List of Sample Output	show ripng general-statistics on page 982
Output Fields	Table 146 on page 982 lists the output fields for the show ripng general-statistics command. Output fields are listed in the approximate order in which they appear.

Table 146: show ripng general-statistics Output Fields

Field Name	Field Description
bad msgs	Number of invalid messages received.
no recv intf	Number of packets received with no matching interface.
curr memory	Amount of memory currently used by RIPng.
max memory	Most memory used by RIPng.

Sample Output

show ripng general-statistics

```

user@host> show ripng general-statistics
RIPng I/O info:
  bad msgs      :      0
  no recv intf  :      0
  curr memory   :      0
  max memory    :      0

```


show ripng neighbor

Syntax	show ripng neighbor <logical-system (all <i>logical-system-name</i>)> < <i>name</i> >
Syntax (EX Series Switch)	show ripng neighbor < <i>name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display information about RIP next-generation (RIPng) neighbors.
Options	<p>none—Display information about all RIPng neighbors.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>name</i>—(Optional) Display detailed information about a specific RIPng neighbor.</p>
Required Privilege Level	view
List of Sample Output	show ripng neighbor on page 984
Output Fields	Table 147 on page 983 lists the output fields for the show ripng neighbor command. Output fields are listed in the approximate order in which they appear.

Table 147: show ripng neighbor Output Fields

Field Name	Field Description
Neighbor	Name of RIPng neighbor.
State	State of the connection: Up or Dn (Down).
Source Address	Source address.
Destination Address	Destination address.
Send	Send options: broadcast , multicast , none , version 1 , or yes .
Recv	Type of packets to accept: both , none , version 1 , or yes .
In Met	Metric added to incoming routes when advertising into RIPng routes that were learned from other protocols.

Sample Output

show ripng neighbor

```
user@host> show ripng neighbor
```

Neighbor	State	Source Address	Dest Address	Send	Recv	In Met
-----	-----	-----	-----	-----	-----	-----
fe-0/0/2.0	Up	fe80::290:69ff:fe68:b002	ff02::9	yes	yes	1

show ripng statistics

Syntax	show ripng statistics <logical-system (all <i>logical-system-name</i>)> < <i>name</i> >
Syntax (EX Series Switch)	show ripng statistics < <i>name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display RIP next generation (RIPng) statistics about messages sent and received on an interface, as well as information received from advertisements from other routing devices.
Options	<p>none—Display RIPng statistics for all neighbors.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p><i>name</i>—(Optional) Display detailed information about a specific RIPng neighbor.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear ripng statistics on page 981
List of Sample Output	show ripng statistics on page 986
Output Fields	Table 148 on page 985 lists the output fields for the show ripng statistics command. Output fields are listed in the approximate order in which they appear.

Table 148: show ripng statistics Output Fields

Field Name	Field Description
RIPng info	<p>Information about RIPng on the specified interface:</p> <ul style="list-style-type: none"> • port—UDP port number used for RIPng. • holddown—Hold-down interval, in seconds. • rts learned—Number of routes learned through RIPng. • rts held down—Number of routes held down by RIPng. • rqsts dropped—Number of received request packets that were dropped. • resps dropped—Number of received response packets that were dropped. • restart—Graceful restart status. Displayed when RIPng is or has been in the process of graceful restart.

Table 148: show ripng statistics Output Fields (*continued*)

Field Name	Field Description
<i>logical-interface</i>	Name of the logical interface and its statistics: <ul style="list-style-type: none"> routes learned—Number of routes learned on the logical interface. routes advertised—Number of routes advertised by the logical interface. timeout—Timeout interval, in seconds. update interval—Interval between routing table updates, in seconds.
Counter	List of counter types: <ul style="list-style-type: none"> Updates Sent—Number of update messages sent. Triggered Updates Sent—Number of triggered update messages sent. Responses Sent—Number of response messages sent. Bad Messages—Number of invalid messages received. Updates Received—Number of RIPng update messages received. Bad Route Entries—Number of RIPng invalid route entry messages received. Updates Ignored—Number of RIPng update messages ignored. RIPng Requests Received—Number of RIPng request messages received. RIPng Requests Ignored—Number of RIPng request messages ignored.
Total	Total number of packets for the selected counter.
Last 5 min	Number of packets for the selected counter in the most recent 5-minute period.
Last minute	Number of packets for the selected counter in the most recent 1-minute period.

Sample Output

show ripng statistics

```

user@host> show ripng statistics
RIPng info: port 521; holddown 120s;
      rts learned  rts held down  rqsts dropped  resps dropped
              0              0              0              0

so-0/1/3.0: 0 routes learned; 1 routes advertised; timeout 180s; update interval
20s
Counter              Total    Last 5 min  Last minute
-----
Updates Sent          934         16          4
Triggered Updates Sent    1          0          0
Responses Sent         0          0          0
Bad Messages          0          0          0
Updates Received        0          0          0
Bad Route Entries       0          0          0
Updates Ignored         0          0          0
RIPng Requests Received  0          0          0
RIPng Requests Ignored   0          0          0

```

CHAPTER 19

Firewall Filter Operational Mode Commands

- `clear firewall`
- `show firewall`
- `show firewall filter version`
- `show firewall log`
- `show firewall prefix-action-stats`
- `show firewall templates-in-use`
- `show policer`

clear firewall

Syntax	clear firewall (all counter <i>counter-name</i> filter <i>filter-name</i> log (all <i>logical-system-name</i>) logical-system <i>logical-system-name</i>)
Syntax (EX Series Switches)	clear firewall (all counter <i>counter-name</i> filter <i>filter-name</i> log (all <i>logical-system-name</i>) policer counter (all counter-id <i>counter-index</i>))
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. logical-system option introduced in Junos OS Release 9.3. log option introduced before Junos OS Release 11.4.
Description	Clear statistics about configured firewall filters.



NOTE: The **clear firewall** command cannot be used to clear the Routing Engine filter counters on a backup Routing Engine that is enabled for graceful Routing Engine switchover (GRES).

If you clear statistics for firewall filters that are applied to Trio-based DPCs and that also use the **prefix-action** action on matched packets, wait at least 5 seconds before you enter the **show firewall prefix-action-stats** command. A 5-second pause between issuing the **clear firewall** and **show firewall prefix-action-stats** commands avoids a possible timeout of the **show firewall prefix-action-stats** command.

Options	all —Clear the packet and byte counts for all filters. On EX Series switches, this option also clears the packet counts for all policer counters. counter <i>counter-name</i> —Clear the packet and byte counts for a filter counter that has been configured with the counter firewall filter action. filter <i>filter-name</i> —Clear the packet and byte counts for the specified firewall filter. log (all <i>logical-system-name</i>) —Clear log entries for IPv4 firewall filters that have then log as an action. Use log all to clear all log entries or log <i>logical-system-name</i> to clear log entries for the specified logical system. logical-system <i>logical-system-name</i> —Clear the packet and byte counts for the specified logical system.
----------------	--

policer counter (all | counter-id *counter-index*)—(EX8200 switches only) Clear all policer counters using the **policer counter all** command, or clear a specific policer counter using the **policer counter counter-id *counter-index*** command. The value of *counter-index* can be 0, 1, or 2.

Required Privilege Level clear

Related Documentation • [show firewall on page 990](#)

List of Sample Output [clear firewall all on page 989](#)
[clear firewall \(counter counter-name\) on page 989](#)
[clear firewall \(filter filter-name\) on page 989](#)
[clear firewall \(policer counter all\) \(EX8200 Switch\) on page 989](#)
[clear firewall \(policer counter counter-id counter-index\) \(EX8200 Switch\) on page 989](#)

Sample Output

clear firewall all

```
user@host> clear firewall all
```

clear firewall (counter counter-name)

```
user@host> clear firewall counter port-filter-counter
```

clear firewall (filter filter-name)

```
user@host> clear firewall filter ingress-port-filter
```

clear firewall (policer counter all) (EX8200 Switch)

```
user@switch> clear firewall policer counter all
```

clear firewall (policer counter counter-id counter-index) (EX8200 Switch)

```
user@switch> clear firewall policer counter counter-id 0
```

show firewall

Syntax	<pre>show firewall <counter <i>counter-name</i>> <detail> <filter <i>filter-name</i>> <log> <logical-system (all <i>logical-system-name</i>)> <terse></pre>
Syntax (EX Series Switches)	<pre>show firewall <counter <i>counter-name</i>> <detail> <filter <i>filter-name</i>> <log <(detail interface <i>interface-name</i>)>> <policer counters <(detail counter-id <i>counter-index</i> <detail>)>> <terse></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Option logical-system introduced in Junos OS Release 9.3.</p> <p>Option terse introduced in Junos OS Release 9.4.</p> <p>Option policer counters introduced in Junos OS Release 12.2 for EX Series switches.</p> <p>Option detail introduced in Junos OS Release 12.3 for EX Series switches.</p> <p>Option detail introduced in Junos OS Release 14.1 for MX Series routers.</p>
Description	Display enhanced statistics and counters for all configured firewall filters.
Options	<p>none—(Optional) Display statistics and counters for all configured firewall filters and counters. For EX Series switches, this command also displays statistics about all configured policers.</p> <p>counter <i>counter-name</i>—(Optional) Name of a filter counter.</p> <p>detail—(EX Series switches and MX Series routers only) (Optional) Display firewall filter statistics and enhanced policer statistics and counters.</p> <p>filter <i>filter-name</i>—(Optional) Name of a configured filter.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>log—(Optional) Display log entries for firewall filters.</p> <p>log <(detail interface <i>interface-name</i>)>—(EX Series switches only) (Optional) Display detailed log entries of firewall activity or log information about a specific interface.</p> <p>policer counters <(detail counter-id <i>counter-index</i> <detail>)>—(EX8200 switches only) (Optional) Display policer counter statistics in brief or in detail.</p> <p>terse—(Optional) Display firewall filter names only.</p>

Required Privilege Level view

- Related Documentation**
- [clear firewall on page 988](#)
 - [show firewall log on page 998](#)
 - *Verifying That Firewall Filters Are Operational*
 - *Verifying That Policers Are Operational*
 - [show policer on page 1005](#)
 - *Enhanced Policer Statistics Overview*
 - *enhanced-policer*

List of Sample Output

[show firewall filter \(MX Series Router and EX Series Switch\) on page 994](#)
[show firewall filter \(non MX Series Router and EX Series Switch\) on page 994](#)
[show firewall filter \(Dynamic Input Filter\) on page 994](#)
[show firewall \(Logical Systems\) on page 994](#)
[show firewall \(counter counter-name\) on page 995](#)
[show firewall log on page 995](#)
[show firewall policer counters \(EX8200 Switch\) on page 995](#)
[show firewall policer counters \(detail\) \(EX8200 Switch\) on page 995](#)
[show firewall policer counters \(counter-id counter-index\) \(EX8200 Switch\) on page 996](#)
[show firewall policer counters \(counter-id counter-index detail\) \(EX8200 Switch\) on page 996](#)
[show firewall detail on page 996](#)

Output Fields [Table 149 on page 991](#) lists the output fields for the **show firewall** command. Output fields are listed in the approximate order in which they appear.

Table 149: show firewall Output Fields

Field Name	Field Description
Filter	<p>Name of a filter that has been configured with the filter statement at the [edit firewall] hierarchy level.</p> <p>Except on EX Series switches:</p> <ul style="list-style-type: none"> • When an interface-specific filter is displayed, the name of the filter is followed by the full interface name and by either -i for an input filter or -o for an output filter. • When dynamic filters are displayed, the name of the filter is followed by the full interface name and by either -in for an input filter or -out for an output filter. When a logical system-specific filter is displayed, the name of the filter is prefixed with two underscore (__) characters and the name of the logical system (for example, __ls1/filter1). • When a service filter is displayed that uses a service set, the separator between the service-set name and the service-filter name is a semicolon (:). <p>NOTE: For bridge family filter, the ip-protocol match criteria is supported only for IPv4 and not for IPv6. This is applicable for line cards that support the Junos Trio chipset, such as the MX 3D MPC line cards.</p>

Table 149: show firewall Output Fields (*continued*)

Field Name	Field Description
Counters	<p>Display filter counter information:</p> <ul style="list-style-type: none"> • Name—Name of a filter counter that has been configured with the counter firewall filter action. • Bytes—Number of bytes that match the filter term under which the counter action is specified. • Packets—Number of packets that matched the filter term under which the counter action is specified. <p>NOTE: On M and T series routers, firewall filters cannot count ip-options packets on a per option type and per interface basis. A limited work around is to use the show pfe statistics ip options command to see ip-options statistics on a per Packet Forwarding Engine (PFE) basis. See <i>show pfe statistics ip</i> for sample output.</p>
Policers	<p>Display policer information:</p> <ul style="list-style-type: none"> • Name—Name of policer. • Bytes—(For two-color policers on MX Series routers and EX Series switches, and for hierarchical policers on MS-DPC, MIC, and MPC interfaces on MX Series routers) Number of bytes that match the filter term under which the policer action is specified. This is only the number out-of-specification (out-of-spec) byte counts, not all the bytes in all packets policed by the policer. For other combinations of policer type, device, and line card type, this field is blank. • Packets—Number of packets that matched the filter term under which the policer action is specified. This is only the number of out-of-specification (out-of-spec) packet counts, not all packets policed by the policer.
Policer Counter Index	(EX8200 switch only) Global management counter ID. The counter ID value (<i>counter-index</i>) can be 0, 1, or 2.
Green	(EX8200 switch only) Number of packets within the limits. The number of packets is smaller than the committed information rate (CIR).
Yellow	(EX8200 switch only) Number of packets partially within the limits. The number of packets is greater than the CIR, but the burst size is within the excess burst size (EBS) limit.
Discard	(EX8200 switch only) Number of discarded packets.
Bytes	(EX8200 switch only) Number of green, yellow, red, or discarded packets in bytes.
Packets	(EX8200 switch only) Number of green, yellow, red, or discarded packets.
Filter name	(EX8200 switch only) Name of the filter with a term associated to a policer.
Term name	(EX8200 switch only) Name of the term associated with a policer.
Policer name	(EX8200 switch only) Name of the policer that is associated with a global management counter.

Table 149: show firewall Output Fields (*continued*)

Field Name	Field Description
PI-t1	<ul style="list-style-type: none">• OOS packet statistics for packets that are marked out-of-specification (out-of-spec) by the policer. Changes to all packets that have out-of-spec actions, such as discard, color marking, or forwarding-class, are included in this counter.• Offered packet statistics for traffic subjected to policing.• Transmitted packet statistics for traffic that is not discarded by the policer. When the policer action is discard, the statistics are the same as the in-spec statistics; when the policer action is non-discard (loss-priority or forwarding-class), the statistics are included in this counter.

Sample Output

show firewall filter (MX Series Router and EX Series Switch)

```
user@host> show firewall filter test
Filter: test
Counters:
Name          Bytes      Packets
Counter-1      0           0
Counter-2      0           0
Policers:
Name          Bytes      Packets
Policer-1     2770       70
```

show firewall filter (non MX Series Router and EX Series Switch)

```
user@host> show firewall filter test
Filter: test
Counters:
Name          Bytes      Packets
Counter-1      0           0
Counter-2      0           0
Policers:
Name          Bytes      Packets
Policer-1     70
```

show firewall filter (Dynamic Input Filter)

```
user@host> show firewall filter dfwd-ge-5/0/0.1-in
Filter: dfwd-ge-5/0/0.1-in
Counters:
Name          Bytes      Packets
c1-ge-5/0/0.1-in 0           0
```

show firewall (Logical Systems)

```
user@host> show firewall

Filter: __lr1/test
Counters:
Name          Bytes      Packets
icmp          420         5
Filter: __default_bpdu_filter__
Filter: __lr1/inet_filter1
Counters:
Name          Bytes      Packets
inet_tcp_count 0           0
inet_udp_count 0           0
Filter: __lr1/inet_filter2
Counters:
Name          Bytes      Packets
inet_icmp_count 0           0
inet_pim_count 0           0
Filter: __lr2/inet_filter1
Counters:
Name          Bytes      Packets
inet_tcp_count 0           0
inet_udp_count 0           0
```

show firewall (counter counter-name)

```

user@host> show firewall counter icmp-counter
Filter: ingress-port-voip-class-filter
Counters:
Name                                     Bytes      Packets
icmp-counter                             0           0

```

show firewall log

```

user@host> show firewall log
Log :

Time      Filter  Action Interface  Protocol  Src Addr
      Dest Addr
08:00:53  pfe      R    ge-1/0/1.0  ICMP      192.168.3.5
      192.168.3.4
08:00:52  pfe      R    ge-1/0/1.0  ICMP      192.168.3.5
      192.168.3.4
08:00:51  pfe      R    ge-1/0/1.0  ICMP      192.168.3.5
      192.168.3.4
08:00:50  pfe      R    ge-1/0/1.0  ICMP      192.168.3.5
      192.168.3.4
08:00:49  pfe      R    ge-1/0/1.0  ICMP      192.168.3.5
      192.168.3.4
08:00:48  pfe      R    ge-1/0/1.0  ICMP      192.168.3.5
      192.168.3.4
08:00:47  pfe      R    ge-1/0/1.0  ICMP      192.168.3.5
      192.168.3.4

```

show firewall policer counters (EX8200 Switch)

```

user@switch> show firewall policer counters
Policer Counter Index 0:

Bytes      Packets
Green:      73      15914
Yellow:     9      1962
Discard:    119    25942

Policer Counter Index 1:

Bytes      Packets
Green:      0      0
Yellow:     0      0
Discard:    0      0

Policer Counter Index 2:

Bytes      Packets
Green:      0      0
Yellow:     0      0
Discard:    0      0

```

show firewall policer counters (detail) (EX8200 Switch)

```

user@switch> show firewall policer counters detail
Policer Counter Index 0:

Bytes      Packets
Green:      73      15914
Yellow:     9      1962
Discard:    119    25942

```

Filter name	Term name	Policer name
myfilter	polcr-term-1	myfilter-polcr-1
inet-filter-ae	ae-snmp	policer-1
inet-filter-ae	ae-ssh	policer-2

Policer Counter Index 1:

	Bytes	Packets
Green:	0	0
Yellow:	0	0
Discard:	0	0

Filter name	Term name	Policer name
-------------	-----------	--------------

Policer Counter Index 2:

	Bytes	Packets
Green:	0	0
Yellow:	0	0
Discard:	0	0

Filter name	Term name	Policer name
-------------	-----------	--------------

show firewall policer counters (counter-id counter-index) (EX8200 Switch)

user@switch> show firewall policer counters counter-id 0

Policer Counter Index 0:

	Bytes	Packets
Green:	73	15914
Yellow:	9	1962
Discard:	119	25942

show firewall policer counters (counter-id counter-index detail) (EX8200 Switch)

user@switch> show firewall policer counters counter-id 0 detail

Policer Counter Index 0:

	Bytes	Packets
Green:	73	15914
Yellow:	9	1962
Discard:	119	25942

Filter name	Term name	Policer name
myfilter	polcr-term-1	myfilter-polcr-1
inet-filter-ae	ae-snmp	policer-1
inet-filter-ae	ae-ssh	policer-2

show firewall detail

user@host> show firewall detail

Filter: __default_bpdu_filter__

Filter: foo

Counters:

Name	Bytes	Packets
c1	17652140	160474

Policers:

Name	Bytes	Packets
P1-t1		
OOS	0	18286
Offered	0	18446744073709376546
Transmitted	0	18446744073709358260

show firewall filter version

Syntax	show firewall filter version <filter-name>
Release Information	Command introduced in Junos OS Release 10.2R2.
Description	Display the version number of the installed firewall filter in the Routing Engine.
Options	<p>none—(Optional) Display the version number of all installed firewall filters.</p> <p>filter-name—(Optional) Name of a configured filter. If you specify the name of a filter, only the version number of that filter is displayed.</p>
Additional Information	The initial version number is 1. This number increments by one when you modify the firewall filter settings or an associated prefix action. The maximum version number is 4,294,967,295. When the version number reaches 4,294,967,295, this number is reset to 1.
Required Privilege Level	view
List of Sample Output	show firewall filter version on page 997
Output Fields	Table 150 on page 997 lists the output fields for the show firewall filter version command. Output fields are listed in the approximate order in which they appear.

Table 150: show firewall filter version Output Fields

Field Name	Field Description
Filter	Name of a filter that has been configured with the filter statement at the [edit firewall] hierarchy level.
Version	Display the version number of the firewall filter.

Sample Output

show firewall filter version

```

user@host> show firewall filter version
Filter version information :
Filter                                     Version
test                                     10

```

show firewall log

Syntax	show firewall log <detail> <interface <i>interface-name</i> > <logical-system (<i>logical-system-name</i> all)>
Syntax (EX Series Switches)	show firewall log <detail> <interface <i>interface-name</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. logical-system option introduced in Junos OS Release 9.3.
Description	Display log information about firewall filters.
Options	none —Display log information about firewall filters. detail —(Optional) Display detailed information. interface <i>interface-name</i> —(Optional) Display log information about a specific interface. logical-system (<i>logical-system-name</i> all) —(Optional) Perform this operation on all logical systems or on a particular system.
Required Privilege Level	view
List of Sample Output	show firewall log on page 999 show firewall log detail on page 999
Output Fields	Table 151 on page 998 lists the output fields for the show firewall log command. Output fields are listed in the approximate order in which they appear.

Table 151: show firewall log Output Fields

Field Name	Field Description
Time of Log	Time that the event occurred.
Filter	<ul style="list-style-type: none"> Displays the name of a configured firewall filter or service filter only if the packet hit the filter's log action in a kernel filter (in the control plane). For any traffic that reaches the Routing Engine, the packets hit the log action in the kernel. For all other logged packets (packet hit the filter's log action in the Packet Forwarding Engine), this field displays pfe instead of a configured filter name.

Table 151: show firewall log Output Fields (*continued*)

Field Name	Field Description
Filter Action	Filter action: <ul style="list-style-type: none"> • A—Accept • D—Discard • R—Reject
Name of Interface	<ul style="list-style-type: none"> • Displays a physical interface name if the packet arrived at a port on a line card. • Displays local if the packet was generated by the device's internal Ethernet interface, em1 or fxp1, which connects the Routing Engine with the router's packet-forwarding components.
Name of protocol	Packet's protocol name: egp , gre , icmp , ipip , ospf , pim , rsvp , tcp , or udp .
Packet length	Length of the packet.
Source address	Packet's source address.
Destination address	Packet's destination address and port.

Sample Output

show firewall log

```

user@host>show firewall log
Time      Filter  Action Interface    Protocol  Src Addr    Dest Addr
13:10:12  pfe      D      rlsq0.902     ICMP     180.1.177.2 180.1.177.1
13:10:11  pfe      D      rlsq0.902     ICMP     180.1.177.2 180.1.177.1

```

show firewall log detail

```

user@host> show firewall log detail
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0Name of protocol: TCP, Packet Length: 50824, Source address:
172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 1020, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0

```

Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
....

show firewall prefix-action-stats

Syntax	show firewall prefix-action-stats filter <i>filter-name</i> prefix-action <i>prefix-action-name</i> <from <i>number</i> to <i>number</i> > <logical-system (<i>logical-system-name</i> all)>
Release Information	Command introduced before Junos OS Release 7.4. logical-system option introduced in Junos OS Release 9.3.
Description	Display prefix action statistics about configured firewall filters. If you clear statistics for firewall filters that are applied to Trio-based MPCs and that also use the prefix-action action on matched packets, wait at least 5 seconds before you enter the show firewall prefix-action-stats command. A 5-second pause between issuing the clear firewall and show firewall prefix-action-stats commands avoids a possible timeout of the show firewall prefix-action-stats command.
Options	filter <i>filter-name</i> —Name of a filter. prefix-action <i>prefix-action-name</i> —Name of a prefix action. from <i>number</i> to <i>number</i> —(Optional) Starting and ending counter or policer. logical-system (<i>logical-system-name</i> all) —(Optional) Perform this operation on all logical systems or on a particular system.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> clear firewall on page 988
List of Sample Output	show firewall prefix-action-stats on page 1001
Output Fields	Table 152 on page 1001 lists the output fields for the show firewall prefix-action-stats command. Output fields are listed in the approximate order in which they appear.

Table 152: show firewall prefix-action-stats Output Fields

Field Name	Field Description
Filter	Filter name. Filters configured for logical systems include the name of the filter prefixed with the two underscore characters (__) and the name of the logical system (for example, __ls1/filter1).

Sample Output

show firewall prefix-action-stats

```
user@host> show firewall prefix-action-stats filter test prefix-action act1
Filter: __ls2/test
```


show firewall templates-in-use

Syntax	show firewall templates-in-use
Release Information	Command introduced in Junos OS Release 12.3.
Description	Display the names of configured filter templates that are currently in use by dynamic subscribers and the number of times each template is referenced.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• clear firewall on page 988• show firewall log on page 998• <i>Verifying That Firewall Filters Are Operational</i>
List of Sample Output	show firewall templates-in-use on page 1004
Output Fields	Table 153 on page 1003 lists the output fields for the show firewall templates-in-use command. Output fields are listed in the approximate order in which they appear.

Table 153: show firewall templates-in-use Output Fields

Field Name	Field Description
Filter Template	Name of a filter that has been configured using the filter statement at either the [edit firewall] or [edit dynamic-profiles <i>profile-name</i> firewall] hierarchy and is being used as a template for dynamic subscriber filtering.
Reference Count	Number of times the filter has been referenced by subscribers accessing the network.

Sample Output

show firewall templates-in-use

```
user@host> show firewall templates-in-use
```

Filter Template	Dynamic Subscribers	Reference Counts
-----		-----
egressFilter		10
ingressFilter		10
dfilter		5
dfilter-pol		5

show policer

Syntax	show policer <detail> <policer-name>
Release Information	Command introduced before Junos OS Release 7.4. Option detail introduced in Junos OS Release 12.3.
Description	Display the number of policed packets for a given policer or an aggregate policer. An aggregate policer is an aggregate of different policers on the same logical interface.
Options	none —Display the number of policed packets for all configured policers. detail —(Optional) Display enhanced statistics and counters for policers. policer-name —(Optional) Display the number of policed packets for the specified policer.
Required Privilege Level	view
List of Sample Output	show policer (MX Series) on page 1006 show policer (non MX Series Router) on page 1006 show policer (Aggregate Policar, non MX Series Router) on page 1006 show policer detail on page 1007
Output Fields	Table 154 on page 1005 lists the output fields for the show policer command. Output fields are listed in the approximate order in which they appear.

Table 154: show policer Output Fields

Field Name	Field Description
Name	Name of the policer.
Bytes	<ul style="list-style-type: none"> (For two-color policers on MX Series routers, and for hierarchical policers on MS-DPC, MIC, and MPC interfaces on MX Series routers)—Total number of bytes policed by the specified policer. For other combinations of policer type, device, and line card type, this field is blank. (T Series and M10i)—Not applicable. The Bytes information is not displayed.
Packets	Total number of packets policed by the specified policer.

Table 154: show policer Output Fields (*continued*)

Field Name	Field Description
Policer detail	<ul style="list-style-type: none"> OOS packet statistics for packets that are marked out-of-specification by the policer. Changes to all packets that have out-of-specification actions, such as discard, color marking, or forwarding-class, are included in this counter. Offered packet statistics for traffic subjected to policing. Transmitted packet statistics for traffic that is not discarded by the policer. When the policer action is discard, the statistics are the same as the within-specification statistics; when the policer action is non-discard (loss-priority or forwarding-class), the statistics are included in this counter.

Sample Output

show policer (MX Series)

```

user@host> show policer
Policers:
Name                                     Bytes          Packets
__default_arp_policer__                 314520         5242
pol-2M-ge-1/2/0.1-inet-i                10372300      103723
pol-2M-ge-1/2/0.1-inet6-i               7727800       77278
pol-2M-ge-1/2/0.1-mp1s-i                7070336       67984
pol-2M-ge-1/2/0.1001-vpls-i             65153700      651537
pol-2M-ge-1/2/0.2001-vpls-i             65180900      651809
pol-2M-ge-1/2/0.3001-ccc-i              62202144      647939

```

show policer (non MX Series Router)

```

user@host> show policer
Policers:
Name                                     Bytes          Packets
__default_arp_policer__                 NA             5242
pol-2M-ge-1/2/0.1-inet-i                NA             103723
pol-2M-ge-1/2/0.1-inet6-i               NA             77278
pol-2M-ge-1/2/0.1-mp1s-i                NA             67984
pol-2M-ge-1/2/0.1001-vpls-i             NA             651537
pol-2M-ge-1/2/0.2001-vpls-i             NA             651809
pol-2M-ge-1/2/0.3001-ccc-i              NA             647939

```

show policer (Aggregate Policar, non MX Series Router)

```

user@host> show policer
Policers:
Name                                     Bytes          Packets
__default_arp_policer__                 NA             0
P1-ae0.0-log_int-o                      NA             0
P2-ge-7/0/2.0-inet-o                    NA             0
P2-ge-7/0/2.0-inet6-o                   NA             0
__policer_tmpl__-term                    NA             0
__policer_tmpl__-fc0                     NA             0
__policer_tmpl__-fc0                     NA             0
__policer_tmpl__-fc1                     NA             0
__policer_tmpl__-fc0                     NA             0
__policer_tmpl__-fc1                     NA             0

```


__policer_tmpl__-fc2	NA	0
__policer_tmpl__-fc0	NA	0
__policer_tmpl__-fc1	NA	0
__policer_tmpl__-fc2	NA	0
__policer_tmpl__-fc3	NA	0

show policer detail

```
user@host> show policer detail
```

Policers:

Name	Bytes	Packets
__default_arp_policer__		
OOS	0	0
Offered	0	496
Transmitted	0	496
P1-xe-1/0/0.0-inet-i		
OOS	0	11329
Offered	0	111188
Transmitted	0	99859

CHAPTER 20

Layer 2 Bridging and Switching Operational Mode Commands

- `clear bridge mac-table`
- `clear error bpdu`
- `clear error mac-rewrite`
- `show bridge domain`
- `show bridge flood`
- `show bridge mac-table`
- `show bridge statistics`
- `show l2-learning global-information`
- `show l2-learning global-mac-count`
- `show l2-learning instance`
- `show l2-learning interface`
- `show mac-rewrite interface`

clear bridge mac-table

Syntax	<code>clear bridge mac-table</code> <code><bridge-domain (all <i>bridge-domain-name</i>)></code> <code><instance <i>instance-name</i>></code> <code><interface <i>interface-name</i>></code> <code><learning-vlan id (all-vlan <i>learning-vlan-id</i>)></code> <code><mac-address></code>
Release Information	Command introduced in Junos OS Release 8.4.
Description	(MX Series routers only) Clear learned Layer 2 address information from the media access control (MAC) address table.
Options	<p>none—Clear all learned Layer 2 address information from the MAC address table.</p> <p>bridge-domain (all <i>bridge-domain-name</i>)—(Optional) Clear learned Layer 2 MAC addresses for all bridging domains or for the specified bridging domain.</p> <p>instance <i>instance-name</i>—(Optional) Clear learned Layer 2 MAC addresses for the specified routing instance.</p> <p>interface <i>interface-name</i>—(Optional) Clear learned Layer 2 MAC addresses for the specified interface.</p> <p>learning-vlan-id (all-vlan <i>learning-vlan-id</i>)—(Optional) Clears learned Layer 2 MAC addresses for all VLANs or for the specified VLAN.</p> <p>mac-address—(Optional) Clear the specified learned Layer 2 address from the MAC address table.</p>
Required Privilege Level	clear
List of Sample Output	clear bridge mac-table on page 1010
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear bridge mac-table

```
user@host> clear bridge mac-table
```

clear error bpd

Syntax	<code>clear error bpd</code> <code><interface <i>interface-name</i>></code>
Release Information	Command introduced in Junos OS Release 9.4.
Description	(MX Series routers only) Clear a bridge protocol data unit (BPDU) error condition caused by the detection of a possible bridging loop from Spanning Tree Protocol (STP) operation.
Options	<code>interface <i>interface-name</i></code> —(Optional) Clear the BPDU error condition for the specified interface.
Required Privilege Level	clear
List of Sample Output	clear error bpd interface on page 1011
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear error bpd interface

```
user@host> clear error bpd interface ge-1/1/1
```

clear error mac-rewrite

Syntax	<code>clear error mac-rewrite</code> <code><interface <i>interface-name</i>></code>
Release Information	Command introduced in Junos OS Release 9.1.
Description	(MX Series routers only) Clear a MAC rewrite error condition caused by the reception of tunneled Cisco Discovery Protocol (CDP), Spanning Tree Protocol (STP), or VLAN Trunk Protocol (VTP) packets on an interface with Layer 2 protocol tunneling enabled.
Options	<code>interface <i>interface-name</i></code> —(Optional) Clear the MAC rewrite error condition for the specified interface.
Required Privilege Level	clear
List of Sample Output	clear error mac-rewrite interface on page 1012
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear error mac-rewrite interface

```
user@host> clear error mac-rewrite interface ge-1/0/1
```

show bridge domain

Syntax	<pre>show bridge domain <brief detail extensive> <bridge-domain (all domain-name)> <instance instance-name> <operational></pre>
Release Information	Command introduced in Junos OS Release 8.4.
Description	(MX Series routers only) Display bridge domain information.
Options	<p>none—Display information for all bridge domains.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>bridge-domain (all domain-name)— (Optional) Display information about all bridge domains or the specified bridge domain.</p> <p>instance instance-name—(Optional) Display information for the specified routing instance.</p> <p>operational—(Optional) Display information for the operational routing instances.</p>
Required Privilege Level	view
List of Sample Output	show bridge domain on page 1013 show bridge domain brief on page 1013 show bridge domain detail on page 1013

Sample Output

show bridge domain

```
user@host> show bridge domain
Instance      Primary Table  Bridging Domain  Type      Active
vs1           bridge.0      vlan100          bridge    2
vs1           bridge.0      vlan200          bridge    0
```

show bridge domain brief

```
user@host> show bridge domain brief
Instance      Primary Table  Bridging Domain  Type      Active
vs1           bridge.0      vlan100          bridge    2
vs1           bridge.0      vlan200          bridge    0
```

show bridge domain detail

```
user@host> show bridge domain detail
```

```
Routing Instance:vs1
  Bridging Domain:vlan100
  Router ID: 0.0.0.0
  Type: bridge                      State: Active
  Interfaces:
    ge-11/0/3.0
    ge-11/1/4.100
    ge-11/1/1.100
    ge-11/1/0.100
    xe-10/2/0.100
    xe-10/0/0.100
  Tables:
    bridge.0                        : 2 macs (2 active)
Routing Instance:vs1
  Bridging Domain:vlan200
  Router ID: 0.0.0.0
  Type: bridge                      State: Active
  Interfaces:
    ge-11/1/0.200
    ge-11/1/1.200
    ge-11/1/4.200
    xe-10/0/0.200
    xe-10/2/0.200
  Tables:
    bridge.0                        : 0 macs (0 active)
```


show bridge flood

Syntax	<pre>show bridge flood <brief detail extensive> <bridge-domain <i>domain-name</i>> <event-queue> <instance <i>instance-name</i>> <route (all-ce-flood all ve-flood alt-root-flood bd-flood mlp-flood re-flood)></pre>
Release Information	Command introduced in Junos OS Release 8.4.
Description	(MX Series routers only) Display bridging flooding information.
Options	<p>none—Display all bridging flooding information for all bridging domains.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>bridge-domain <i>domain-name</i>—(Optional) Display bridging flooding information for the specified bridge domain.</p> <p>event-queue—(Optional) Display the queue of pending bridge flood events.</p> <p>instance <i>instance-name</i>—(Optional) Display bridging flooding information for the specified routing instance.</p> <p>route (all-ce-flood all ve-flood alt-root-flood bd-flood mlp-flood re-flood)—(Optional) Display the following:</p> <ul style="list-style-type: none"> all-ce-flood—Display the route for flooding traffic to all customer edge routers if no-local-switching is enabled. all-ve-flood—Display the route for flooding traffic to all VPLS edge routers if no-local-switching is enabled. alt-root-flood—Display the Spanning Tree Protocol (STP) alt-root flooding route used for the interface. bd-flood—Display the route for flooding traffic of a bridge domain if no-local-switching is not enabled. mlp-flood—Display the route for flooding traffic to MAC learning chips. re-flood—Display the route for Routing Engine flooding to all interfaces.
Required Privilege Level	view
List of Sample Output	show bridge flood on page 1016 show bridge flood brief on page 1016 show bridge flood detail on page 1016 show bridge flood extensive on page 1017
Output Fields	to be provided

Sample Output

show bridge flood

```

user@host> show bridge flood
Name: __juniper_private1__
CEs: 0
VEs: 0
Flood Routes:
  Prefix   Type      Owner                NhType  NhIndex
  0x36/16  MLP_FLOOD __vs1+vlan100__      flood   426
  0x3a/16  MLP_FLOOD __vs1+vlan200__      flood   428
Name: vs1::vlan100
CEs: 6
VEs: 0
Flood Routes:
  Prefix   Type      Owner                NhType  NhIndex
  0x35/16  ALL_FLOOD __vs1+vlan100__      flood   425
  0x35/16  RE_FLOOD  __vs1+vlan100__      flood   425
  0x3780/17 ALT_ROOT_RT ge-11/0/3.0          flood   425
  0x3b80/17 ALT_ROOT_RT ge-11/1/4.100         flood   425
  0x3c80/17 ALT_ROOT_RT ge-11/1/1.100         flood   425
  0x3d80/17 ALT_ROOT_RT ge-11/1/0.100         flood   425
  0x3e80/17 ALT_ROOT_RT xe-10/2/0.100         flood   425
  0x3f80/17 ALT_ROOT_RT xe-10/0/0.100         flood   425
Name: vs1::vlan200
CEs: 5
VEs: 0
Flood Routes:
  Prefix   Type      Owner                NhType  NhIndex
  0x39/16  ALL_FLOOD __vs1+vlan200__      flood   427
  0x39/16  RE_FLOOD  __vs1+vlan200__      flood   427
  0x4180/17 ALT_ROOT_RT ge-11/1/0.200         flood   427
  0x4080/17 ALT_ROOT_RT ge-11/1/1.200         flood   427
  0x4280/17 ALT_ROOT_RT ge-11/1/4.200         flood   427
  0x4480/17 ALT_ROOT_RT xe-10/0/0.200         flood   427
  0x4380/17 ALT_ROOT_RT xe-10/2/0.200         flood   427

```

show bridge flood brief

```

user@host> show bridge flood brief
Name                Active CEs    Active VEs
__juniper_private1__ 0              0
vs1::vlan100         6              0
vs1::vlan200         5              0

```

show bridge flood detail

```

user@host> show bridge flood detail
Name: __juniper_private1__
CEs: 0
VEs: 0
Flood Routes:
  Prefix   Type      Owner                NhType  NhIndex
  0x36/16  MLP_FLOOD __vs1+vlan100__      flood   426
  0x3a/16  MLP_FLOOD __vs1+vlan200__      flood   428
Name: vs1::vlan100
CEs: 6
VEs: 0
Flood Routes:
  Prefix   Type      Owner                NhType  NhIndex

```

```

0x35/16  ALL_FLOOD    __vs1+vlan100__  flood    425
0x35/16  RE_FLOOD     __vs1+vlan100__  flood    425
0x3780/17 ALT_ROOT_RT ge-11/0/3.0      flood    425
0x3b80/17 ALT_ROOT_RT ge-11/1/4.100    flood    425
0x3c80/17 ALT_ROOT_RT ge-11/1/1.100    flood    425
0x3d80/17 ALT_ROOT_RT ge-11/1/0.100    flood    425
0x3e80/17 ALT_ROOT_RT xe-10/2/0.100     flood    425
0x3f80/17 ALT_ROOT_RT xe-10/0/0.100     flood    425
Name: vs1::vlan200
CEs: 5
VEs: 0
Flood Routes:
  Prefix  Type      Owner          NhType  NhIndex
  0x39/16  ALL_FLOOD __vs1+vlan200__ flood    427
  0x39/16  RE_FLOOD  __vs1+vlan200__ flood    427
  0x4180/17 ALT_ROOT_RT ge-11/1/0.200    flood    427
  0x4080/17 ALT_ROOT_RT ge-11/1/1.200    flood    427
  0x4280/17 ALT_ROOT_RT ge-11/1/4.200    flood    427
  0x4480/17 ALT_ROOT_RT xe-10/0/0.200     flood    427
  0x4380/17 ALT_ROOT_RT xe-10/2/0.200     flood    427

```

show bridge flood extensive

```

user@host> show bridge flood extensive
Name: __juniper_private1__
CEs: 0
VEs: 0
Flood route prefix: 0x36/16
Flood route type: MLP_FLOOD
Flood route owner: __vs1+vlan100__
Nexthop type: flood
Nexthop index: 426
  Interfaces Flooding to:
    Name          Type      NhType      Index
    1c-11/0/0.32769 LC
    1c-10/2/0.32769 LC
    1c-10/0/0.32769 LC
    1c-11/1/0.32769 LC

Flood route prefix: 0x3a/16
Flood route type: MLP_FLOOD
Flood route owner: __vs1+vlan200__
Nexthop type: flood
Nexthop index: 428
  Interfaces Flooding to:
    Name          Type      NhType      Index
    1c-10/0/0.32769 LC
    1c-10/2/0.32769 LC
    1c-11/1/0.32769 LC
Name: vs1::vlan100
CEs: 6
VEs: 0

Flood route prefix: 0x35/16
Flood route type: ALL_FLOOD
Flood route owner: __vs1+vlan100__
Nexthop type: flood
Nexthop index: 425
  Interfaces Flooding to:
    Name          Type      NhType      Index
    ge-11/0/3.0    CE

```

```
ge-11/1/4.100    CE
ge-11/1/1.100    CE
ge-11/1/0.100    CE
xe-10/2/0.100    CE
xe-10/0/0.100    CE
```

```
Flood route prefix: 0x35/16
Flood route type: RE_FLOOD
Flood route owner: __vs1+vlan100__
Nexthop type: flood
Nexthop index: 425
```

```
  Interfaces Flooding to:
```

Name	Type	NhType	Index
ge-11/0/3.0	CE		
ge-11/1/4.100	CE		
ge-11/1/1.100	CE		
ge-11/1/0.100	CE		
xe-10/2/0.100	CE		
xe-10/0/0.100	CE		

```
Flood route prefix: 0x3780/17
Flood route type: ALT_ROOT_RT
Flood route owner: ge-11/0/3.0
Nexthop type: flood
Nexthop index: 425
```

```
  Interfaces Flooding to:
```

Name	Type	NhType	Index
ge-11/0/3.0	CE		
ge-11/1/4.100	CE		
ge-11/1/1.100	CE		
ge-11/1/0.100	CE		
xe-10/2/0.100	CE		
xe-10/0/0.100	CE		

```
Flood route prefix: 0x3b80/17
Flood route type: ALT_ROOT_RT
Flood route owner: ge-11/1/4.100
Nexthop type: flood
Nexthop index: 425
```

```
  Interfaces Flooding to:
```

Name	Type	NhType	Index
ge-11/0/3.0	CE		
ge-11/1/4.100	CE		
ge-11/1/1.100	CE		
ge-11/1/0.100	CE		
xe-10/2/0.100	CE		
xe-10/0/0.100	CE		

```
Flood route prefix: 0x3c80/17
Flood route type: ALT_ROOT_RT
Flood route owner: ge-11/1/1.100
Nexthop type: flood
Nexthop index: 425
```

```
  Interfaces Flooding to:
```

Name	Type	NhType	Index
ge-11/0/3.0	CE		
ge-11/1/4.100	CE		
ge-11/1/1.100	CE		
ge-11/1/0.100	CE		
xe-10/2/0.100	CE		
xe-10/0/0.100	CE		

Flood route prefix: 0x3d80/17
 Flood route type: ALT_ROOT_RT
 Flood route owner: ge-11/1/0.100
 Nexthop type: flood
 Nexthop index: 425

Interfaces Flooding to:

Name	Type	NhType	Index
ge-11/0/3.0	CE		
ge-11/1/4.100	CE		
ge-11/1/1.100	CE		
ge-11/1/0.100	CE		
xe-10/2/0.100	CE		
xe-10/0/0.100	CE		

Flood route prefix: 0x3e80/17
 Flood route type: ALT_ROOT_RT
 Flood route owner: xe-10/2/0.100
 Nexthop type: flood
 Nexthop index: 425

Interfaces Flooding to:

Name	Type	NhType	Index
ge-11/0/3.0	CE		
ge-11/1/4.100	CE		
ge-11/1/1.100	CE		
ge-11/1/0.100	CE		
xe-10/2/0.100	CE		
xe-10/0/0.100	CE		

Flood route prefix: 0x3f80/17
 Flood route type: ALT_ROOT_RT
 Flood route owner: xe-10/0/0.100
 Nexthop type: flood
 Nexthop index: 425

Interfaces Flooding to:

Name	Type	NhType	Index
ge-11/0/3.0	CE		
ge-11/1/4.100	CE		
ge-11/1/1.100	CE		
ge-11/1/0.100	CE		
xe-10/2/0.100	CE		
xe-10/0/0.100	CE		

Name: vs1::vlan200

CEs: 5

VEs: 0

Flood route prefix: 0x39/16
 Flood route type: ALL_FLOOD
 Flood route owner: __vs1+vlan200__
 Nexthop type: flood
 Nexthop index: 427

Interfaces Flooding to:

Name	Type	NhType	Index
ge-11/1/0.200	CE		
ge-11/1/1.200	CE		
ge-11/1/4.200	CE		
xe-10/0/0.200	CE		
xe-10/2/0.200	CE		

Flood route prefix: 0x39/16
 Flood route type: RE_FLOOD

Flood route owner: __vs1+vlan200__

Nexthop type: flood

Nexthop index: 427

Interfaces Flooding to:

Name	Type	NhType	Index
ge-11/1/0.200	CE		
ge-11/1/1.200	CE		
ge-11/1/4.200	CE		
xe-10/0/0.200	CE		
xe-10/2/0.200	CE		

Flood route prefix: 0x4180/17

Flood route type: ALT_ROOT_RT

Flood route owner: ge-11/1/0.200

Nexthop type: flood

Nexthop index: 427

Interfaces Flooding to:

Name	Type	NhType	Index
ge-11/1/0.200	CE		
ge-11/1/1.200	CE		
ge-11/1/4.200	CE		
xe-10/0/0.200	CE		
xe-10/2/0.200	CE		

Flood route prefix: 0x4080/17

Flood route type: ALT_ROOT_RT

Flood route owner: ge-11/1/1.200

Nexthop type: flood

Nexthop index: 427

Interfaces Flooding to:

Name	Type	NhType	Index
ge-11/1/0.200	CE		
ge-11/1/1.200	CE		
ge-11/1/4.200	CE		
xe-10/0/0.200	CE		
xe-10/2/0.200	CE		

Flood route prefix: 0x4280/17

Flood route type: ALT_ROOT_RT

Flood route owner: ge-11/1/4.200

Nexthop type: flood

Nexthop index: 427

Interfaces Flooding to:

Name	Type	NhType	Index
ge-11/1/0.200	CE		
ge-11/1/1.200	CE		
ge-11/1/4.200	CE		
xe-10/0/0.200	CE		
xe-10/2/0.200	CE		

Flood route prefix: 0x4480/17

Flood route type: ALT_ROOT_RT

Flood route owner: xe-10/0/0.200

Nexthop type: flood

Nexthop index: 427

Interfaces Flooding to:

Name	Type	NhType	Index
ge-11/1/0.200	CE		
ge-11/1/1.200	CE		
ge-11/1/4.200	CE		
xe-10/0/0.200	CE		

xe-10/2/0.200 CE

Flood route prefix: 0x4380/17

Flood route type: ALT_ROOT_RT

Flood route owner: xe-10/2/0.200

Nexthop type: flood

Nexthop index: 427

Interfaces Flooding to:

Name	Type	NhType	Index
ge-11/1/0.200	CE		
ge-11/1/1.200	CE		
ge-11/1/4.200	CE		
xe-10/0/0.200	CE		
xe-10/2/0.200	CE		

show bridge mac-table

Syntax	<code>show bridge mac-table</code> <code><brief count detail extensive></code> <code><bridge-domain (all <i>bridge-domain-name</i>)></code> <code><global-count></code> <code><interface <i>interface-name</i>></code> <code><mac-address></code> <code><vlan-id (all-vlan <i>vlan-id</i>)></code>
Release Information	Command introduced in Junos OS Release 8.4.
Description	(MX Series routers only) Display Layer 2 MAC address information.
Options	<p>none—Display all learned Layer 2 MAC address information.</p> <p>brief count detail extensive—(Optional) Display the specified level of output.</p> <p>bridge-domain (all <i>bridge-domain-name</i>)—(Optional) Display learned Layer 2 MAC addresses for all bridging domains or for the specified bridging domain.</p> <p>global-count—(Optional) Display the total number of learned Layer 2 MAC addresses on the system.</p> <p>instance <i>instance-name</i>—(Optional) Display learned Layer 2 MAC addresses for the specified routing instance.</p> <p>interface <i>interface-name</i>—(Optional) Display learned Layer 2 MAC addresses for the specified interface.</p> <p>mac-address—(Optional) Display the specified learned Layer 2 MAC address information.</p> <p>vlan-id (all-vlan <i>vlan-id</i>)—(Optional) Display learned Layer 2 MAC addresses for all VLANs or for the specified VLAN.</p>
Additional Information	When Layer 2 protocol tunneling is enabled, the tunneling MAC address 01:00:0c:cd:cd:d0 is installed in the MAC table. When the Cisco Discovery Protocol (CDP), Spanning Tree Protocol (STP), or VLAN Trunk Protocol (VTP) is configured for Layer 2 protocol tunneling on an interface, the corresponding protocol MAC address is installed in the MAC table.
Required Privilege Level	view
List of Sample Output	show bridge mac-table on page 1023 show bridge mac-table (with VXLAN enabled) on page 1024 show bridge mac-table count on page 1024 show bridge mac-table detail on page 1025
Output Fields	Table 155 on page 1023 describes the output fields for the show bridge mac-table command. Output fields are listed in the approximate order in which they appear.

Table 155: show bridge mac-table Output fields

Field Name	Field Description
Routing instance	Name of the routing instance.
Bridging domain	Name of the bridging domain.
MAC address	MAC address or addresses learned on a logical interface.
MAC flags	Status of MAC address learning properties for each interface: <ul style="list-style-type: none"> • S—Static MAC address is configured. • D—Dynamic MAC address is configured. • L—Locally learned MAC address is configured. • C—Control MAC address is configured. • SE—MAC accounting is enabled. • NM—Non-configured MAC. • R—Remote PE MAC address is configured.
Logical interface	Name of the logical interface.
MAC count	Number of MAC addresses learned on the specific routing instance or interface.
Learning interface	Name of the logical interface on which the MAC address was learned.
Learning VLAN	VLAN ID of the routing instance or bridge domain in which the MAC address was learned.
VXLAN ID/VXLAN	VXLAN Network Identifier (VNI)
Layer 2 flags	Debugging flags signifying that the MAC address is present in various lists.
Epoch	Spanning Tree Protocol epoch number identifying when the MAC address was learned. Used for debugging.
Sequence number	Sequence number assigned to this MAC address. Used for debugging.
Learning mask	Mask of the Packet Forwarding Engines where this MAC address was learned. Used for debugging.
IPC generation	Creation time of the logical interface when this MAC address was learned. Used for debugging.

Sample Output

show bridge mac-table

```

user@host> show bridge mac-table
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned, C -Control MAC
          SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

```

```

Routing instance : default-switch
Bridging domain : test1, VLAN : 1
MAC          MAC      Logical   NH      RTR
address      flags    interface Index   ID
01:00:0c:cc:cc:cc S,NM    NULL
01:00:0c:cc:cc:cd S,NM    NULL
01:00:0c:cd:cd:d0 S,NM    NULL
64:87:88:6a:17:d0 D        ae0.1
64:87:88:6a:17:f0 D        ae0.1

```

show bridge mac-table (with VXLAN enabled)

```

user@host> show bridge mac-table
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned
          SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

```

```

Routing instance : default-switch
Bridging domain : vlan-1, VLAN : 1
VXLAN: Id : 100, Multicast group: 226.1.1.1
MAC          MAC      Logical   NH      RTR
address      flags    interface Index   ID
00:01:01:00:01:f7 D,SE    vtep.1052010
00:03:00:32:01:f7 D,SE    vtep.1052011
00:00:21:11:11:10 DL        ge-1/0/0.0
00:00:21:11:11:11 DL        ge-1/1/0.0

```

```

Routing instance : default-switch
Bridging domain : vlan-2, VLAN : 2, VXLAN : 200
VXLAN: Id : 200, Multicast group: 226.1.1.2
MAC          MAC      Logical   NH      RTR
address      flags    interface Index   ID
00:02:01:33:01:f7 D,SE    vtep.1052010
00:04:00:14:01:f7 D,SE    vtep.1052011
00:00:21:11:21:10 DL        ge-1/0/0.1
00:00:21:11:21:11 DL        ge-1/1/0.1

```

show bridge mac-table count

```

user@host> show bridge mac-table count
2 MAC address learned in routing instance vs1 bridge domain vlan100

```

MAC address count per interface within routing instance:

Logical interface	MAC count
ge-11/0/3.0	1
ge-11/1/4.100	0
ge-11/1/1.100	0
ge-11/1/0.100	0
xe-10/2/0.100	1
xe-10/0/0.100	0

MAC address count per learn VLAN within routing instance:

Learn VLAN ID	MAC count
0	2

```

0 MAC address learned in routing instance vs1 bridge domain vlan200

```

MAC address count per interface within routing instance:

Logical interface	MAC count
-------------------	-----------

ge-11/1/0.200	0
ge-11/1/1.200	0
ge-11/1/4.200	0
xe-10/0/0.200	0
xe-10/2/0.200	0

MAC address count per learn VLAN within routing instance:

Learn VLAN ID	MAC count
0	0

show bridge mac-table detail

```
user@host> show bridge mac-table detail
```

MAC address: 00:00:00:19:1c:db

Routing instance: vs1

Bridging domain: vlan100

Learning interface: ge-11/0/3.0 Learning VLAN: 0

Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel

Epoch: 4 Sequence number: 0

Learning mask: 0x800 IPC generation: 0

MAC address: 00:00:00:59:3a:2f

Routing instance: vs1

Bridging domain: vlan100

Learning interface: xe-10/2/0.100 Learning VLAN: 0

Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel

Epoch: 7 Sequence number: 0

Learning mask: 0x400 IPC generation: 0

show bridge statistics

Syntax	show bridge statistics <bridge-domain <i>domain-name</i>> <instance <i>instance-name</i>>
Release Information	Command introduced in Junos OS Release 8.4.
Description	(MX Series routers only) Display bridge statistics.
Options	none —Display bridge statistics for all bridge domains in all routing instances. bridge-domain <i>domain-name</i> —(Optional) Display statistics for the specified bridge domain. instance <i>instance-name</i> —(Optional) Display statistics for the specified routing instance.
Required Privilege Level	view
List of Sample Output	show bridge statistics on page 1026

Sample Output

show bridge statistics

```
user@host> show bridge statistics
Information for routing instance:

Routing instance : __juniper_private1__
  Index: 1                               Sequence number: 0
  MAC limit: 5000                         MACs learned: 0
  Static MACs learned: 0                  Non config Static MACs learned: 0
  Handle: 0x829e800

Information for routing instance:

Routing instance : vs1
  Bridging domain : vlan100
  Index: 3                               Sequence number: 0
  MAC limit: 5120                         MACs learned: 2
  Static MACs learned: 0                  Non config Static MACs learned: 0
  Handle: 0x829e400
  Flags: Bridge instance, Config defined, VLAN : 100
  Local interface: ge-11/0/3.0, Index: 79
    Broadcast packets: 1
    Broadcast bytes : 65
    Multicast packets: 0
    Multicast bytes : 0
    Flooded packets : 0
    Flooded bytes : 0
    Unicast packets : 358624489
    Unicast bytes : 23310592305
    Current MAC count: 1 (Limit 1024)
  Local interface: ge-11/1/4.100, Index: 84
    Broadcast packets: 0
    Broadcast bytes : 0
    Multicast packets: 0
    Multicast bytes : 0
```

```

Flooded packets : 0
Flooded bytes : 0
Unicast packets : 0
Unicast bytes : 0
Current MAC count: 0 (Limit 1024)
Local interface: ge-11/1/1.100, Index: 86
Broadcast packets: 0
Broadcast bytes : 0
Multicast packets: 0
Multicast bytes : 0
Flooded packets : 0
Flooded bytes : 0
Unicast packets : 0
Unicast bytes : 0
Current MAC count: 0 (Limit 1024)
Local interface: ge-11/1/0.100, Index: 87
Broadcast packets: 0
Broadcast bytes : 0
Multicast packets: 0
Multicast bytes : 0
Flooded packets : 0
Flooded bytes : 0
Unicast packets : 0
Unicast bytes : 0
Current MAC count: 0 (Limit 1024)
Local interface: xe-10/2/0.100, Index: 88
Broadcast packets: 0
Broadcast bytes : 0
Multicast packets: 0
Multicast bytes : 0
Flooded packets : 0
Flooded bytes : 0
Unicast packets : 358627393
Unicast bytes : 23310781065
Current MAC count: 1 (Limit 1024)
Local interface: xe-10/0/0.100, Index: 89
Broadcast packets: 0
Broadcast bytes : 0
Multicast packets: 0
Multicast bytes : 0
Flooded packets : 0
Flooded bytes : 0
Unicast packets : 0
Unicast bytes : 0
Current MAC count: 0 (Limit 1024)

```

Information for routing instance:

```

Routing instance : vs1
Bridging domain : vlan200
Index: 4                      Sequence number: 0
MAC limit: 5120               MACs learned: 0
Static MACs learned: 0       Non config Static MACs learned: 0
Handle: 0x829e600
Flags: Bridge instance, Config defined, VLAN : 200
Local interface: ge-11/1/0.200, Index: 90
Broadcast packets: 0
Broadcast bytes : 0
Multicast packets: 0
Multicast bytes : 0
Flooded packets : 0

```

```
Flooded bytes      : 0
Unicast packets    : 0
Unicast bytes      : 0
Current MAC count: 0 (Limit 1024)
Local interface: ge-11/1/1.200, Index: 91
Broadcast packets: 0
Broadcast bytes    : 0
Multicast packets: 0
Multicast bytes    : 0
Flooded packets    : 0
Flooded bytes      : 0
Unicast packets    : 0
Unicast bytes      : 0
Current MAC count: 0 (Limit 1024)
Local interface: ge-11/1/4.200, Index: 92
Broadcast packets: 0
Broadcast bytes    : 0
Multicast packets: 0
Multicast bytes    : 0
Flooded packets    : 0
Flooded bytes      : 0
Unicast packets    : 0
Unicast bytes      : 0
Current MAC count: 0 (Limit 1024)
Local interface: xe-10/0/0.200, Index: 93
Broadcast packets: 0
Broadcast bytes    : 0
Multicast packets: 0
Multicast bytes    : 0
Flooded packets    : 0
Flooded bytes      : 0
Unicast packets    : 0
Unicast bytes      : 0
Current MAC count: 0 (Limit 1024)
Local interface: xe-10/2/0.200, Index: 94
Broadcast packets: 4
Broadcast bytes    : 260
Multicast packets: 0
Multicast bytes    : 0
Flooded packets    : 0
Flooded bytes      : 0
Unicast packets    : 0
Unicast bytes      : 0
Current MAC count: 0 (Limit 1024)
```

show l2-learning global-information

Syntax	show l2-learning global-information
Release Information	Command introduced in Junos OS Release 8.4.
Description	(MX Series routers only) Display Layer 2 learning process-related information for the entire router.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show l2-learning global-information on page 1029
Output Fields	Table 156 on page 1029 describes the output fields for the show l2-learning global-information command. Output fields are listed in the approximate order in which they appear.

Table 156: show l2-learning global-information Output Fields

Field Name	Field Description
MAC aging interval	Configured timeout interval, in seconds, for all MAC table entries.
MAC learning	Status of MAC learning: Enabled or Disabled .
MAC statistics	Status of MAC accounting: Enabled or Disabled .
MAC limit Count	Configured maximum limit on the number of MAC addresses that can be learned.
MAC limit hit flag	Status of the learned MAC limit hit flag: Enabled (the learned MAC exceeds the global MAC limit) or Disabled (the learned MAC does not exceed the global MAC limit).
MAC packet action drop	Status of action to drop packets after the configured MAC address limit is reached: Enabled (packets are dropped) or Disabled (packets are forwarded).

Sample Output

show l2-learning global-information

```

user@host> show l2-learning global-information
Global Configuration:

MAC aging interval      : 300
MAC learning            : Enabled
MAC statistics          : Disabled
MAC limit Count         : 393215

```

MAC limit hit flag : Disabled
MAC packet action drop: Disabled

show l2-learning global-mac-count

Syntax	show l2-learning global-mac-count
Release Information	Command introduced in Junos OS Release 9.3.
Description	(MX Series routers only) Display the total number of dynamic and static MAC addresses learned for the entire router.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show l2-learning global-mac-count on page 1031
Output Fields	Displays the total number of dynamic and static MAC addresses learned for the entire router.

Sample Output

show l2-learning global-mac-count

```
user@host> show l2-learning global-mac-count
100 dynamic and static MAC addresses learned globally
```

show l2-learning instance

Syntax	show l2-learning instance
Release Information	(MX Series routers only) Command introduced in Junos OS Release 8.4.
Description	Display Layer 2 learning properties for all the configured routing instances.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show l2-learning instance on page 1032
Output Fields	Table 157 on page 1032 describes the output fields for the show l2-learning instance command. Output fields are listed in the approximate order in which they appear.

Table 157: show l2-learning instance Output Fields

Field Name	Field Description
Routing Instance	Name of routing instance.
Bridging Domain	Name of bridging domain. On MX Series routers you can use the show l2-learning instance <extensive> command option to display the Bridge Service-id information which includes the Config Service ID and the Active Service ID.
Index	Number associated with the routing instance or bridging domain.
Logical System	Name of logical system or Default if no logical system is configured.
Routing instance flags	Status of Layer 2 learning properties for each routing instance: <ul style="list-style-type: none"> • DL—MAC learning is disabled. • SE—MAC accounting is enabled. • AD—Packets are dropped after MAC address limit is reached. • LH—The maximum number of MAC addresses has been learned on the routing instance. The routing instance is not able to learn any additional MAC addresses.
MAC limit	Maximum number of MAC addresses that can be learned from each interface in the routing instance or bridging domain.

Sample Output

show l2-learning instance

```

user@host> show l2-learning instance
Information for routing instance:

Routing Instance flags (DL -disable learning, SE -stats enabled,
```

AD -packet action drop, LH -mac limit hit)

Routing Instance	Bridging Domain	Index	Logical System	Routing flags	MAC limit
__juniper_private1__		1	Default		5000
vs1	vlan100	3	Default		5120
vs1	vlan200	4	Default		5120

show l2-learning interface

Syntax	show l2-learning interface
Release Information	Command introduced in Junos OS Release 8.4.
Description	(MX Series routers only) Display Layer 2 learning information for all the interfaces.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show l2-learning interface on page 1034
Output Fields	Table 158 on page 1034 describes the output fields for the show l2-learning interface command. Output fields are listed in the approximate order in which they appear.

Table 158: show l2-learning interface Output Fields

Field Name	Field Description
Logical interface	Name of the logical interface.
Index	Index of the interface.
Routing Instance	Number of the routing instance to which the interface belongs.
Interface device	Value of the order in which the Junos OS finds and initializes the interface.
Logical interface flags	Status of Layer 2 learning properties for each interface: <ul style="list-style-type: none"> • DL—MAC learning is disabled. • SE—MAC accounting is enabled. • AD—Packets are dropped after the MAC interface limit is reached. • MAC limit—Maximum number of MAC addresses that can be learned from the interface.

Sample Output

show l2-learning interface

```
user@host> show l2-learning interface
Information for interface family:
```

```
Logical Interface flags (DL -disable learning, SE -stats enabled,
AD -packet action drop, LH -mac limit hit)
```

Logical interface	Index	Routing instance	Interface device	Logical Interface flags	MAC limit
ge-11/0/3.0	79	3	136		1024
ge-11/1/4.100	84	3	150		1024
ge-11/1/1.100	86	3	147		1024

ge-11/1/0.100	87	3	146	1024
xe-10/2/0.100	88	3	144	1024
xe-10/0/0.100	89	3	129	1024
ge-11/1/0.200	90	4	146	1024
ge-11/1/1.200	91	4	147	1024
ge-11/1/4.200	92	4	150	1024
xe-10/0/0.200	93	4	129	1024
xe-10/2/0.200	94	4	144	1024

show mac-rewrite interface

Syntax	show mac-rewrite interface <brief detail> <interface-name>
Release Information	Command introduced in Junos OS Release 9.1.
Description	(MX Series routers only) Display Layer 2 protocol tunneling information.
Options	brief detail —(Optional) Display the specified level of output. interface <i>interface-name</i> —(Optional) Display Layer 2 protocol tunneling information for the specified interface.
Required Privilege Level	view
List of Sample Output	show mac-rewrite interface on page 1036
Output Fields	Table 159 on page 1036 lists the output fields for the show mac-rewrite interface command. Output fields are listed in the approximate order in which they appear.

Table 159: show mac-rewrite interface Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface that has Layer 2 protocol tunneling configured on it.	brief detail
Protocols	Layer 2 protocols being tunneled on this interface: Cisco Discovery Protocol (CDP), Spanning Tree Protocol (STP), Per-VLAN Spanning Tree Plus (PVSTP+), or VLAN Trunk Protocol (VTP)	brief detail

Sample Output

show mac-rewrite interface

```

user@host> show mac-rewrite interface
Interface      Protocols
ge-1/0/5      STP VTP CDP PVSTP+

```

CHAPTER 21

VPN Operational Mode Commands

- `clear vpls mac-address`
- `clear vpls mac-table`
- `request l2circuit-switchover`
- `show dynamic-tunnels database`
- `show hfrr profiles`
- `show ingress-replication mvpn`
- `show l2circuit connections`
- `show l2vpn connections`
- `show mvpn c-multicast`
- `show mvpn instance`
- `show mvpn neighbor`
- `show vpls connections`
- `show vpls flood event-queue`
- `show vpls flood instance`
- `show vpls flood route`
- `show vpls mac-table`
- `show vpls statistics`

clear vpls mac-address

Syntax	<code>clear vpls mac-address</code> <code><instance <i>instance-name</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code> <code><mac-address></code>
Release Information	Command introduced before Junos OS Release 7.4.
Description	(T Series and M Series routers, except for the M160 router) Clear media access control (MAC) address entries from the virtual private LAN service (VPLS) table.
Options	none —Clear all MAC address entries from the VPLS table for all routing instances. instance <i>instance-name</i> —(Optional) Clear all MAC address entries for a VPLS instance from the VPLS table. logical-system (all <i>logical-system-name</i>) —(Optional) Perform this operation on all logical systems or on a particular logical system. mac-address —(Optional) Clear a specific MAC address in a VPLS instance from the VPLS table.
Required Privilege Level	maintenance
List of Sample Output	clear vpls mac-address on page 1038
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear vpls mac-address

```
user@host> clear vpls mac-address
```


clear vpls mac-table

Syntax	clear vpls mac-table <instance <i>instance-name</i> > <interface <i>interface-name</i> > <logical-system (all <i>logical-system-name</i>)> <mac-address> <vlan-id>
Release Information	Command introduced before Junos OS Release 9.5.
Description	(MX Series routers) Clear media access control (MAC) addresses from the virtual private LAN service (VPLS) MAC table.
Options	<p>none—Clear all MAC addresses from the VPLS table for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Clear all MAC addresses for a VPLS instance from the VPLS table.</p> <p>interface <i>interface-name</i>—(Optional) Clear all MAC addresses for a VPLS interface from the VPLS table.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>mac-address—(Optional) Clear a specific MAC address in a VPLS instance from the VPLS table.</p> <p>vlan-id—(Optional) Clear MAC addresses on a specified VLAN (0 through 4095).</p>
Required Privilege Level	maintenance
List of Sample Output	clear vpls mac-table on page 1039
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear vpls mac-table

```
user@host> clear vpls mac-table
```

request l2circuit-switchover

Syntax	<code>request l2circuit-switchover</code> <code><logical-system (all logical-system-name)></code> <code><neighbor <i>address</i>></code> <code><virtual-circuit-id <i>identifier</i>></code>
Release Information	Command introduced in Junos OS Release 9.2.
Description	Manually trigger a switch from the active pseudowire to the redundant pseudowire. This command can be useful when performing network maintenance.
Options	<p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>neighbor <i>address</i>—(Optional) Trigger a switch of all of the active pseudowire connections with the specified neighbor to their respective redundant pseudowires.</p> <p>virtual-circuit-id <i>identifier</i>—(Optional) Trigger a switch from the active pseudowire connection of the specified Layer 2 circuit to its redundant pseudowire.</p>
Required Privilege Level	maintenance
List of Sample Output	request l2circuit-switchover virtual-circuit-id on page 1040
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request l2circuit-switchover virtual-circuit-id

```
user@host>request l2circuit-switchover virtual-circuit-id 12
```

show dynamic-tunnels database

Syntax	show dynamic-tunnels database <destination> <logical-system (all <i>logical-system-name</i>)> <table <i>routing-table-name</i> >
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display dynamic tunnel database information.
Options	<p>none—Display dynamic tunnel database information for all destinations and routing tables.</p> <p>destination—(Optional) Display database entries for the specified IP address (with optional destination prefix length) only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>table <i>routing-table-name</i>—(Optional) Display database entries for the specified table only.</p>
Required Privilege Level	view
List of Sample Output	show dynamic-tunnels database (Tunnel Is Up) on page 1042 show dynamic-tunnels database (No Tunnel PIC) on page 1042 show dynamic-tunnels database (Tunnel Is Expiring) on page 1042 show dynamic-tunnels database (Destination Specified) on page 1042
Output Fields	Table 160 on page 1041 lists the output fields for the show dynamic-tunnels database command. Output fields are listed in the approximate order in which they appear.

Table 160: show dynamic-tunnels database Output Fields

Field Name	Field Description
Table	Name of the routing table (for example, inet.0).
Destination-network	Destination IP address and subnet.
Tunnel to	Destination IP address and prefix of the tunnel.
State	State of the tunnel: Up , Up (expires in <i>nn:nn:nn</i>seconds) , or Dn (down).
Reference count	Number of routes across the dynamic tunnel that are currently being resolved.
Next-hop type	Type of tunnel: GRE.
Source address	Source IP address of the tunnel.

Table 160: show dynamic-tunnels database Output Fields (*continued*)

Field Name	Field Description
Next-hop	IP address of the destination interface.
State	State of the destination interface: Up, Dn, or Dn (no tunnel pic).

Sample Output

show dynamic-tunnels database (Tunnel Is Up)

```

user@host> show dynamic-tunnels database
Table: inet.3

Destination-network: 10.255.120.94/32
Tunnel to: 10.255.120.94/32 State: Up
Reference count: 2
Next-hop type: gre
Source address: 10.255.120.92
Next hop: gr-4/3/0.32769
State: Up

```

show dynamic-tunnels database (No Tunnel PIC)

```

user@host> show dynamic-tunnels database
Table: inet.3

Destination-network: 10.255.120.94/32
Tunnel to: 10.255.120.94/32 State: Dn
Reference count: 2
Next-hop type: gre
Source address: 10.255.120.92
State: Dn (no tunnel pic)

```

show dynamic-tunnels database (Tunnel Is Expiring)

```

user@host> show dynamic-tunnels database
Table: inet.3

Destination-network: 10.255.120.94/32
Tunnel to: 10.255.120.94/32 State: Up (expires in 00:14:56 seconds)
Reference count: 0
Next-hop type: gre
Source address: 10.255.120.92
Next hop: gr-4/3/0.32769
State: Up

```

show dynamic-tunnels database (Destination Specified)

```

user@host> show dynamic-tunnels database 10.255.120.94
Table: inet.3

Destination-network: 10.255.120.94/32
Tunnel to: 10.255.120.94/32 State: Up
Reference count: 2
Next-hop type: gre
Source address: 10.255.120.92

```

Next hop: gr-4/3/0.32769
State: Up

show hfr profiles

Syntax	show hfr profiles <brief extensive>
Release Information	Command introduced in Junos OS Release 12.2.
Description	<p>Display host fast reroute (HFRR) profile information.</p> <p>HFRR adds a precomputed protection path into the Packet Forwarding Engine, such that if a link between a provider edge device and a server farm becomes unusable for forwarding, the Packet Forwarding Engine can use another path without having to wait for the router or the protocols to provide updated forwarding information.</p>
Options	<p>none—Display information about HRFF profiles.</p> <p>brief extensive—(Optional) Display the specified level of output.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Host Fast Reroute</i>
List of Sample Output	show hfr profiles on page 1045
Output Fields	<p>Table 161 on page 1044 describes the output fields for the show hfr profiles command. Output fields are listed in the approximate order in which they appear.</p>

Table 161: show hfr profiles Output Fields

Field Name	Field Description
HFRR pointer	
HFRR current state	Status of the HFRR profile: HFRR_ACTIVE , HFRR_INACTIVE , HFRR_IFLH-NOT-CONF , and so on.
HFRR Prefix limit blackout timer expiry (in secs)	Time interval between an HFRR profile becoming inactive on exceeding the ARP prefix limit, and the profile starting the SYNC process.
HFRR prefix limit hit count	Number of times that an HFRR profile becomes inactive on exceeding the ARP prefix limit.
HFRR protected IFL name	Interface configured for the HFRR feature.
HFRR protected IFL handle	
HFRR routing instance name	The routing instance in which the HFRR interface is configured.
HFRR routing instance handle	

Table 161: show hfrf profiles Output Fields (*continued*)

Field Name	Field Description
HFRR sync BG scheduled	
HFRR RTS filter on	
HFRR delete BG scheduled	
HFRR ARP prefix limit	Configured ARP prefix limit.
HFRR ARP supplementary blackout timeout (in mins)	Supplementary time-out value configured for profile to be inactive when it hits ARP prefix limit.
HFRR number of ARP routes learned	Number of ARP routes learned on the configured interface.
HFRR number of FRR routes created	Number of ARP routes created on the configured interface.

Sample Output

show hfrf profiles

```

user@host> show hfrf profiles
HFRR pointer: 0x9254000
HFRR current state: HFRR_ACTIVE
HFRR Prefix limit blackout timer expiry (in secs): 0
HFRR prefix limit hit count: 0
HFRR protected IFL name: ge-4/1/0.0
HFRR protected IFL handle: 0x9248738
HFRR routing instance name: test
HFRR routing instance handle: 0x9145740
HFRR sync BG scheduled: NO
HFRR RTS filter on: YES
HFRR delete BG scheduled: NO
HFRR ARP prefix limit: 0
HFRR ARP supplementary blackout timeout (in mins): 1
HFRR number of ARP routes learned: 4
HFRR number of FRR routes created: 2

```

show ingress-replication mvpn

Syntax	show ingress-replication mvpn
Release Information	Command introduced in Junos OS Release 10.4.
Description	Display the state and configuration of the ingress replication tunnels created for the MVPN application when using the mpls-internet-multicast routing instance type.
Required Privilege Level	View
List of Sample Output	show ingress-replication mvpn on page 1046
Output Fields	Table 162 on page 1046 lists the output fields for the show ingress-replication mvpn command. Output fields are listed in the approximate order in which they appear.

Table 162: show ingress-replication mvpn

Field Name	Field Description
Ingress tunnel	Identifies the MVPN ingress replication tunnel.
Application	Identifies the application (MVPN).
Unicast tunnels	List of unicast tunnels in use.
Leaf address	Address of the tunnel.
Tunnel type	Identifies the unicast tunnel type.
Mode	Indicates whether the tunnel was created as a new tunnel for the ingress replication, or if an existing tunnel was used.
State	Indicates whether the tunnel is Up or Down.

Sample Output

show ingress-replication mvpn

```

user@host> show ingress-replication mvpn
Ingress Tunnel: mvpn:1
  Application: MVPN
  Unicast tunnels
    Leaf Address    Tunnel-type    Mode    State
    10.255.245.2    P2P LSP       New     Up
    10.255.245.4    P2P LSP       New     Up
Ingress Tunnel: mvpn:2
  Application: MVPN
  Unicast tunnels
    Leaf Address    Tunnel-type    Mode    State
    10.255.245.2    P2P LSP       Existing Up

```


show l2circuit connections

Syntax	<code>show l2circuit connections</code> <code><brief extensive summary></code> <code><down up up-down></code> <code><history></code> <code><interface <i>interface-name</i>></code> <code><logical-system (all <i>logical-system-name</i>)></code> <code><neighbor <i>neighbor</i>></code> <code><status></code>
Release Information	Command introduced before Junos OS Release 7.4. Display enhancements in Junos OS Release 9.6. Display enhancements in Junos OS Release 10.2. Display enhancements in Junos OS Release 12.1. Display enhancements in Junos OS Release 13.2.
Description	Display status information about Layer 2 virtual circuits from the local provider edge (PE) router to its neighbors.
Options	<p>none—Display standard information about Layer 2 virtual circuits on all interfaces for all neighbors.</p> <p>brief extensive summary—(Optional) Display the specified level of output. Use history to display information about connection history. Use status to display information about the connection and interface status.</p> <p>down up up-down—(Optional) Display nonoperational, operational, or both kinds of connections.</p> <p>history—(Optional) Display information about connection history.</p> <p>interface <i>interface-name</i>—(Optional) Show all Layer 2 virtual circuits on an interface.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>neighbor <i>neighbor</i>—(Optional) IP address of a specific neighbor.</p> <p>status—(Optional) Display information about the connection and interface status.</p>
Required Privilege Level	view
List of Sample Output	show l2circuit connections on page 1051 show l2circuit connections interface on page 1052 show l2circuit connections extensive on page 1052 show l2circuit connections extensive (Pseudowire Redundancy with Hot Standby) on page 1053
Output Fields	Table 163 on page 1049 lists the output fields for the show l2circuit connections command. Output fields are listed in the approximate order in which they appear.

Table 163: show l2circuit connections Output Fields

Field Name	Field Description
Layer-2 Circuit Connections	Displays the legends for connection and interface status.
Neighbor	Remote PE neighbor.
Interface	Logical PE-to-CE interface on which the virtual circuit is configured.
Type	VC type: rmt (remote) or loc (local).
Legend for connection status (St)	<p>Status of the virtual circuit connection:</p> <ul style="list-style-type: none"> • EI—The local virtual circuit interface is configured with an encapsulation that is not supported. • MM—The two routers do not agree on an MTU value, which causes an MTU mismatch. • EM—The encapsulation type received on this virtual circuit from the neighbor does not match the local virtual circuit interface encapsulation type. • CM—The two routers do not agree on a control word, which causes a control word mismatch. • VM—The remote and local VLAN IDs do not match across the Layer 2 circuit. • OL—No advertisement has been received for this virtual circuit from the neighbor. There is no outgoing label available for use by this virtual circuit. • NC—The interface is not configured as a CCC or TCC interface. • BK—The virtual circuit has switched to a backup connection. • CB—The remote PE router is advertising a different cell bundle from that configured on the local PE router. • LD—The connection to the local site is signaled down, because the CE-facing interface to the local site is down. • RD—The remote neighbor is down. It has signaled a problem using the pseudowire status code. • NP—The router detects that interface hardware is not present. The hardware may be offline, a PIC may not be of the desired type, or the interface may be configured in a different routing instance. • Dn—The virtual circuit is down. • VC-Dn—The virtual circuit is down because there is no tunnel LSP from the local PE router to the neighbor. • UP—The virtual circuit is operational. • CF—The router cannot find enough bandwidth to the remote router to satisfy the Layer 2 circuit bandwidth requirement. • IB—The bit rate is incompatible for Time Division Multiplexing (TDM). • TDM—TDM is not configured correctly. • ST—The virtual circuit has been switched to a standby connection. • SP—The virtual circuit connection is using a static pseudowire. • RS—The remote site is in a standby state. • XX—The virtual circuit is down for an unknown reason. This is a programming error.
Time last up	Date and time the virtual circuit was last operational.

Table 163: show l2circuit connections Output Fields (*continued*)

Field Name	Field Description
# Up trans	Number of times the virtual circuit came up.
<i>local-interface-name</i>	Name of the local PE-to-CE interface.
Status	Status of the local interface.
Up	Interface is operational.
Dn	Interface is not operational.
NP	Not present. Interface does not exist.
DS	Disabled. Interface has been administratively disabled.
WE	Wrong encapsulation. The interface is not configured as CCC.
UN	Interface status is initialized.
Encapsulation	Encapsulation of the local interface.
Remote PE	Prefix of the remote PE router.
Negotiated control-word	Whether the use of the control word has been negotiated for this virtual circuit: Yes (Null) or No .
Incoming label	Label used by the remote side of the virtual circuit to send packets destined to the local side. This label is routed to the local virtual circuit interface.
Outgoing label	Label used by the local side of the virtual circuit to send packets to the remote side of the virtual circuit. Packets originated on the local virtual circuit interface are encapsulated with this label before being placed on the tunnel LSP to the neighbor for this virtual circuit. This label is allocated by the neighbor and is used in demultiplexing incoming packets destined for this virtual circuit.
Negotiated PW status TLV	Displays the pseudowire status type, length, and value (TLV). TLVs are a method of encoding variable-length or optional information. If the pseudowire status TLV is used, the corresponding local or neighbor PE router status code is also displayed.
local PW status code	If the pseudowire status TLV is used, displays the local PE router status code.
Neighbor PW status code	If the pseudowire status TLV is used, displays the neighbor PE router status code.
Local interface	Name of the local interface used for the Layer 2 circuit connection.
Status	Status of the local interface (Up or Down).

Table 163: show l2circuit connections Output Fields (*continued*)

Field Name	Field Description
Encapsulation	Encapsulation configured for the local interface.
APS-active	Indicates that the interface belongs to the working circuit.
APS-inactive	Indicates that the interface belongs to the protect circuit.
Connection protection	Whether or not connection protection is configured for the Layer 2 circuit to the neighbor: Yes or No .
VC bandwidth	Bandwidth requirement of the Layer 2 circuit.
Time	Time at which the event occurred.
Connection History	<p>Event types logged in history.</p> <ul style="list-style-type: none"> • loc intf up—Local virtual circuit interface went up. • loc intf down—Local virtual circuit interface went down. • In lbl Update—Incoming label has been updated. • Out lbl Update—Outgoing label has been updated. • PE route changed—Route to PE router has been updated. • PE route down—Route to PE router is down. • rmt side marked—Remote side is marked. • VC Dn—Remote side indicated that its end of the virtual circuit is down (if the tunnel LSP from the remote side to the local side is down). • status update timer—Status update timer processing. It computes the state of the virtual circuit, and determines whether it should be advertised to or withdrawn from the remote side.

Sample Output

show l2circuit connections

```
user@host> show l2circuit connections
```

```
Layer-2 Circuit Connections:
```

```
Legend for connection status (St)
```

```

EI -- encapsulation invalid      NP -- interface h/w not present
MM -- mtu mismatch              Dn -- down
EM -- encapsulation mismatch    VC-Dn -- Virtual circuit Down
CM -- control-word mismatch     Up -- operational
VM -- vlan id mismatch          CF -- Call admission control failure
OL -- no outgoing label         IB -- TDM incompatible bitrate
NC -- intf encaps not CCC/TCC   TM -- TDM misconfiguration
BK -- Backup Connection         ST -- Standby Connection
CB -- rcvd cell-bundle size bad SP -- Static Pseudowire
LD -- local site signaled down  RS -- remote site standby
RD -- remote site signaled down HS -- hot standby
XX -- unknown

```

```
Legend for interface status
```

```
Up -- operational
```

```

Dn -- down
Neighbor: 10.255.245.51
  Interface                Type St    Time last up          # Up trans
ge-2/0/2.600(vc 5)        rmt  Up    Dec 7 18:11:18 2009      1
  Remote PE: 10.255.245.51, Negotiated control-word: No
  Incoming label: 299856, Outgoing label: 299808
  Negotiated PW status TLV: No
  Local interface: ge-2/0/2.600, Status: Up, Encapsulation: VLAN

```

Sample Output

show l2circuit connections interface

```

user@host> show l2circuit connections interface t1-2/0/0:1:1.0
Layer-2 Circuit Connections:

Legend for connection status (St)
EI -- encapsulation invalid      NP -- interface h/w not present
MM -- mtu mismatch              Dn -- down
EM -- encapsulation mismatch     VC-Dn -- Virtual circuit Down
CM -- control-word mismatch     Up -- operational
VM -- vlan id mismatch          CF -- Call admission control failure
OL -- no outgoing label         IB -- TDM incompatible bitrate
NC -- intf encaps not CCC/TCC   TM -- TDM misconfiguration
BK -- Backup Connection         ST -- Standby Connection
CB -- rcvd cell-bundle size bad SP -- Static Pseudowire
LD -- local site signaled down  RS -- remote site standby
RD -- remote site signaled down HS -- hot standby
XX -- unknown

Legend for interface status
Up -- operational
Dn -- down
Neighbor: 10.1.1.1
  Interface                Type St    Time last up          # Up trans
t1-2/0/0:1:1.0(vc 1)(SP) rmt  Up    Apr 27 04:21:02 2011      1
  Remote PE: 10.1.1.1, Negotiated control-word: Yes (Non-null)
  Incoming label: 1010001, Outgoing label: 1000001
  Negotiated PW status TLV: No
  Local interface: t1-1/0/0:1:1.0, Status: Up, Encapsulation: SATOP-T1,
APS-active
  Local interface: t1-2/0/0:1:1.0, Status: Up, Encapsulation: SATOP-T1,
APS-inactive

```

Sample Output

show l2circuit connections extensive

```

user@host> show l2circuit connections extensive
Layer-2 Circuit Connections:

Legend for connection status (St)
EI -- encapsulation invalid      NP -- interface h/w not present
MM -- mtu mismatch              Dn -- down
EM -- encapsulation mismatch     VC-Dn -- Virtual circuit Down
CM -- control-word mismatch     Up -- operational
VM -- vlan id mismatch          CF -- Call admission control failure
OL -- no outgoing label         IB -- TDM incompatible bitrate
NC -- intf encaps not CCC/TCC   TM -- TDM misconfiguration
BK -- Backup Connection         ST -- Standby Connection
CB -- rcvd cell-bundle size bad SP -- Static Pseudowire

```

LD -- local site signaled down RS -- remote site standby
 RD -- remote site signaled down HS -- hot standby
 XX -- unknown

Legend for interface status

Up -- operational
 Dn -- down

Neighbor: 10.255.49.149

Interface	Type	St	Time last up	# Up trans
ae0.0(vc 100)	rmt	Up	Aug 31 09:36:12 2009	1
Remote PE: 10.255.49.149, Negotiated control-word: Yes (Null)				
Incoming label: 299824, Outgoing label: 299776				
Negotiated PW status TLV: Yes				
Local PW status code: 0x00000000, Neighbor PW status code: 0x00000000				
Local interface: ae0.0, Status: Up, Encapsulation: ETHERNET				
Connection protection: Yes				
Connection History:				
Aug 31 09:36:12 2009	status update timer			
Aug 31 09:36:12 2009	PE route changed			
Aug 31 09:36:12 2009	Out lbl Update			299776
Aug 31 09:36:12 2009	In lbl Update			299824
Aug 31 09:36:12 2009	loc intf up			ae0.0

Sample Output

show l2circuit connections extensive (Pseudowire Redundancy with Hot Standby)

user@host>show l2circuit connections extensive

Layer-2 Circuit Connections:

Legend for connection status (St)

EI -- encapsulation invalid	NP -- interface h/w not present
MM -- mtu mismatch	Dn -- down
EM -- encapsulation mismatch	VC-Dn -- Virtual circuit Down
CM -- control-word mismatch	Up -- operational
VM -- vlan id mismatch	CF -- Call admission control failure
OL -- no outgoing label	IB -- TDM incompatible bitrate
NC -- intf encaps not CCC/TCC	TM -- TDM misconfiguration
BK -- Backup Connection	ST -- Standby Connection
CB -- rcvd cell-bundle size bad	SP -- Static Pseudowire
LD -- local site signaled down	RS -- remote site standby
RD -- remote site signaled down	HS -- Hot-standby Connection
XX -- unknown	

Legend for interface status

Up -- operational
 Dn -- down

Neighbor: 88.0.0.101

Interface	Type	St	Time last up	# Up trans
ge-1/3/2.600(vc 1)	rmt	Up	Jan 24 11:00:26 2013	1
Remote PE: 88.0.0.101, Negotiated control-word: Yes (Null)				
Incoming label: 299776, Outgoing label: 299776				
Negotiated PW status TLV: Yes				
Local PW status code: 0x00000000, Neighbor PW status code: 0x00000000				
Local interface: ge-1/3/2.600, Status: Up, Encapsulation: VLAN				
Connection History:				
Jan 24 11:00:26 2013	status update timer			
Jan 24 11:00:26 2013	PE route changed			
Jan 24 11:00:26 2013	Out lbl Update			299776
Jan 24 11:00:26 2013	In lbl Update			299776
Jan 24 11:00:26 2013	loc intf up			ge-1/3/2.600

Neighbor: 88.0.0.102

Interface	Type	St	Time last up	# Up trans
ge-1/3/2.600(vc 2)	rmt	HS	-----	----

Remote PE: 88.0.0.102, Negotiated control-word: Yes (Null)

Incoming label: 299792, Outgoing label: 299776

Negotiated PW status TLV: Yes

local PW status code: 0x00000020, Neighbor PW status code: 0x00000000

Local interface: ge-1/3/2.600, Status: Up, Encapsulation: VLAN

show l2vpn connections

Syntax	<pre>show l2vpn connections <brief extensive> <down up up-down> <history> <instance <i>instance</i>> <instance-history> <local-site <i>local-site</i>> <logical-system (all <i>logical-system-name</i>)> <remote-site <i>remote-site</i>> <status> <summary></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>instance-history option introduced in Junos OS Release 12.3R2.</p>
Description	Display Layer 2 virtual private network (VPN) connections.
Options	<p>none—Display all Layer 2 VPN connections for all routing instances.</p> <p>brief extensive—(Optional) Display the specified level of output.</p> <p>down up up-down—(Optional) Display nonoperational, operational, or both kinds of connections.</p> <p>history—(Optional) Display information about connection history.</p> <p>instance <i>instance</i>—(Optional) Display connections for the specified routing instance only.</p> <p>instance-history—(Optional) Display information about connection history for a particular instance.</p> <p>local-site <i>local-site</i>—(Optional) Display connections for the specified Layer 2 VPN local site name or ID only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>remote-site <i>remote-site</i>—(Optional) Display connection for the specified Layer 2 VPN remote site ID only.</p> <p>status—(Optional) Display information about the connection and interface status.</p> <p>summary—(Optional) Display summary of all Layer 2 VPN connections information.</p>
Required Privilege Level	view
List of Sample Output	<p>show l2vpn connections on page 1058</p> <p>show l2vpn connections extensive on page 1059</p> <p>show l2vpn connections extensive (VPWS) on page 1059</p>

Output Fields Table 164 on page 1056 lists the output fields for the **show l2vpn connections** command. Output fields are listed in the approximate order in which they appear.

Table 164: show l2vpn connections Output Fields

Field Name	Field Description
Instance	Name of Layer 2 VPN instance.
L2vpn-id	For BGP autodiscovery, a globally unique Layer 2 VPN community identifier for the instance.
Local-ID	BGP local-address assigned to the local routing device.
Local site	Name of local site.
Local source-attachment-id	For FEC 129, the VPWS source attachment identifier. The point-to-point nature of VPWS requires that you specify the source access individual identifier (SAII) and the target access individual identifier (TAII). This SAII-TAII pair defines a unique pseudowire between two PE devices.
Target-attachment-id	For FEC 129, the VPWS target attachment identifier. If the configured target identifier matches a source identifier advertised by a remote PE device by way of a BGP auto-discovery message, the pseudowire between that source-target pair is signaled. If there is no match between an advertised source identifier and the configured target identifier, the pseudowire is not established.
Interface name	Name of interface.
Remote Site ID	Remote site ID.
Label Offset	Numbers within the label block that are skipped to find the next label base.
Label-base	Advertises the first label in a block of labels. A remote PE router uses this first label when sending traffic toward the advertising PE router.
Range	Advertises the label block size.
status-vector	Bit vector advertising the state of local PE-CE circuits to remote PE routers. A bit value of 0 indicates that the local circuit and LSP tunnel to the remote PE router are up, whereas a value of 1 indicates either one or both are down.
connection-site	Name of the connection site.
Type	Type of connection: loc (local) or rmt (remote).
St	Status of the connection. (For a list of possible values, see the Legend for connection status (St) field.)
Time last up	Time that the connection was last in the Up condition.
# Up trans	Number of transitions from Down to Up condition.

Table 164: show l2vpn connections Output Fields (*continued*)

Field Name	Field Description
Local circuit	Address and status of local circuit.
Remote circuit	Address and status of remote circuit.
St	<p>Status of the Layer 2 VPN connection (corresponds with Legend for Connection Status):</p> <ul style="list-style-type: none"> • EI—The local Layer 2 VPN interface is configured with an encapsulation that is not supported. • EM—The encapsulation type received on this Layer 2 VPN connection from the neighbor does not match the local Layer 2 VPN connection interface encapsulation type. • VC-Dn—The virtual circuit is currently down. • CM—The two routers do not agree on a control word, which causes a control word mismatch. • CN—The virtual circuit is not provisioned properly. • OR—The label associated with the virtual circuit is out of range. • OL—No advertisement has been received for this virtual circuit from the neighbor. There is no outgoing label available for use by this virtual circuit. • LD—All of the CE-facing interfaces to the local site are down. Therefore, the connection to the local site is signaled as down to the other PE routers. No pseudowires can be established. • RD—All the interfaces to the remote neighbor are down. Therefore, the remote site has been signaled as down to the other PE routers. No pseudowires can be established. • LN—The local site has lost path selection to the remote site and therefore no pseudowires can be established from this local site. • RN—The remote site has lost path selection to a local site or other remote site and therefore no pseudowires are established to this remote site. • XX—The Layer 2 VPN connection is down for an unknown reason. This is a programming error. • NC—The interface encapsulation is not configured as an appropriate CCC, TCC, or Layer 2 VPN encapsulation. • WE—The encapsulation configured for the interface does not match the encapsulation configured for the associated connection within the Layer 2 VPN routing instance. • NP—The router detects that interface hardware is not present. The hardware might be offline, a PIC might not be of the desired type, or the interface might be configured in a different routing instance. • -> —Only the outbound connection is up. • <- —Only the inbound connection is up. • Up—The Layer 2 VPN connection is operational. • Dn—The Layer 2 VPN connection is down. • CF—The router cannot find enough bandwidth to the remote router to satisfy the Layer 2 VPN connection bandwidth requirement. • SC—The local site identifier matches the remote site identifier. No pseudowire can be established between these two sites. You should configure different values for the local and remote site identifiers.

Table 164: show l2vpn connections Output Fields (*continued*)

Field Name	Field Description
	<ul style="list-style-type: none"> LM—The local site identifier is not the minimum designated, meaning it is not the lowest. There is another local site with a lower site identifier. Pseudowires are not being established to this local site, and the associated local site identifier is not being used to distribute Layer 2 VPN label blocks. However, this is not an error state. Traffic continues to be forwarded to the PE router interfaces connected to the local sites when the local sites are in this state. RM—The remote site identifier is not the minimum designated, meaning it is not the lowest. There is another remote site connected to the same PE router which has lower site identifier. The PE router cannot establish a pseudowire to this remote site and the associated remote site identifier cannot be used to distribute VPLS label blocks. However, this is not an error state. Traffic can continue to be forwarded to the PE router interface connected to this remote site when the remote site is in this state. IL—The incoming packets for the Layer 2 VPN connection have no MPLS label.
Remote PE	Address of the remote provider edge router.
Incoming label	Name of the incoming label.
Outgoing label	Name of the outgoing label.
Time	Date and time of Layer 2 VPN connection event.
Event	Type of event.
Interface/Lbl/PE	Interface, label, or PE router.

Sample Output

show l2vpn connections

```

user@host> show l2vpn connections
L2VPN Connections :
Instance : vpna
Local site: 2 (ce-2)
offset: 1, range: 3, label-base: 32768
  connection-site      Type St  Time last up      # Up trans
  3 (3)                loc  Up   Jul 18 20:45:46 2001      1
    Local circuit: fe-0/0/0.1, Status: Up
    Remote circuit: fe-0/0/3.0, Status: Up
  1                    rmt  Up   Jul 18 21:47:25 2001      1
    Local circuit: fe-0/0/0.0, Status: Up
    Remote PE: 192.0.2.1
    Incoming label: 32768, Outgoing label: 32769
Local site: 3 (ce-3)
offset: 1, range: 2, label-base: 33792
  connection-site      Type St  Time last up      # Up trans
  2 (ce-b)             loc  Up   Jul 18 20:45:46 2001      1
    Local circuit: fe-0/0/0.1, Status: Up
    Remote circuit: fe-0/0/3.0, Status: Up
  1                    rmt  Up   Jul 18 21:47:25 2001      1

```

```

Local circuit: fe-0/0/3.1, Status: Up
Remote PE: 192.0.2.1
Incoming label: 33792, Outgoing label: 32770

```

show l2vpn connections extensive

```

user@host> show l2vpn connections extensive
L2VPN Connections:

```

Legend for connection status (St)

```

EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down   NP -- interface hardware not present
CM -- control-word mismatch     -> -- only outbound connection is up
CN -- circuit not provisioned   <- -- only inbound connection is up
OR -- out of range             Up -- operational
OL -- no outgoing label        Dn -- down
LD -- local site signaled down  CF -- call admission control failure
RD -- remote site signaled down SC -- local and remote site ID collision
LN -- local site not designated LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unknown connection status IL -- no incoming label

```

Instance: vpn1

Local site: SITE-A (1)

Number of local interfaces: 1

Number of local interfaces up: 1

ge-1/0/0.517

2

Label-base	Offset	Size	Range	Preference
800000	1	2	2	100

status-vector: 0

connection-site	Type	St	Time last up	# Up trans
2	rmt	Up	Mar 5 18:24:30 2014	1

Remote PE: 192.0.2.34, Negotiated control-word: Yes (Null)

Incoming label: 800001, Outgoing label: 800000

Local interface: ge-0/0/1.517, Status: Up, Encapsulation: VLAN

Connection History:

Apr 25 11:18:09 2014	PE route changed	
Apr 25 11:18:09 2014	Out lbl Update	800000
Apr 25 11:18:09 2014	In lbl Update	800001
Apr 25 11:18:09 2014	loc intf up	ge-0/0/1.517

show l2vpn connections extensive (VPWS)

```

user@host> show l2vpn connections
Layer-2 VPN connections:

```

Legend for connection status (St)

```

EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch     WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down   NP -- interface hardware not present
CM -- control-word mismatch     -> -- only outbound connection is up
CN -- circuit not provisioned   <- -- only inbound connection is up
OR -- out of range             Up -- operational
OL -- no outgoing label        Dn -- down
LD -- local site signaled down  CF -- call admission control failure
RD -- remote site signaled down SC -- local and remote site ID collision
LN -- local site not designated LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unknown connection status IL -- no incoming label
MM -- MTU mismatch             MI -- Mesh-Group ID not available

```

BK -- Backup connection ST -- Standby connection
PF -- Profile parse failure PB -- Profile busy
RS -- remote site standby SN -- Static Neighbor
LB -- Local site not best-site RB -- Remote site not best-site
VM -- VLAN ID mismatch

Legend for interface status

Up -- operational
Dn -- down

Instance: FEC129-VPWS

L2vpn-id: 100:100

Number of local interfaces: 1

Number of local interfaces up: 1

ge-2/0/5.0

Local source-attachment-id: 1 (ONE)

Target-attachment-id	Type	St	Time last up	# Up trans
2	rmt	Up	Nov 28 16:16:14 2012	1

Remote PE: 198.51.100.2, Negotiated control-word: No

Incoming label: 299792, Outgoing label: 299792

Local interface: ge-2/0/5.0, Status: Up, Encapsulation: ETHERNET

Connection History:

Nov 28 16:16:14 2012	status update timer	
Nov 28 16:16:14 2012	PE route changed	
Nov 28 16:16:14 2012	Out lbl Update	299792
Nov 28 16:16:14 2012	In lbl Update	299792
Nov 28 16:16:14 2012	loc intf up	ge-2/0/5.0

show mvpn c-multicast

Syntax	show mvpn c-multicast <extensive summary> <instance-name <i>instance-name</i> >
Release Information	Command introduced in Junos OS Release 8.4.
Description	Display the multicast VPN customer multicast route information.
Options	<p>extensive summary—(Optional) Display the specified level of output.</p> <p>instance-name <i>instance-name</i>—(Optional) Display output for the specified routing instance.</p>
Required Privilege Level	view
List of Sample Output	show mvpn c-multicast on page 1062 show mvpn c-multicast summary on page 1062 show mvpn c-multicast extensive on page 1062
Output Fields	Table 165 on page 1061 lists the output fields for the show mvpn c-multicast command. Output fields are listed in the approximate order in which they appear.

Table 165: show mvpn c-multicast Output Fields

Field Name	Field Description	Level of Output
Instance	Name of the VPN routing instance.	summary extensive none
C-mcast IPv4 (S:G)	Customer router IPv4 multicast address.	extensive none
Ptnl	Provider tunnel attributes, <i>tunnel type:tunnel source, tunnel destination group</i> .	extensive none
St	State: <ul style="list-style-type: none"> • DS—Represents (S,G) and is created due to (*,G) • RM—Remote VPN route learned from the remote PE router • St display blank—SSM group join 	extensive none
MVPN instance	Name of the multicast VPN routing instance	extensive none
C-multicast IPv4 route count	Number of customer multicast IPv4 routes associated with the multicast VPN routing instance.	summary
C-multicast IPv6 route count	Number of customer multicast IPv6 routes associated with the multicast VPN routing instance.	summary

Sample Output

show mvpn c-multicast

```
user@host> show mvpn c-multicast
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
  C-mcast IPv4 (S:G)          Ptnl          St
  192.168.195.78/32:225.5.5.5/32 PIM-SM:10.255.14.144, 239.1.1.1      RM
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-B
  C-mcast IPv4 (S:G)          Ptnl          St
  192.168.195.94/32:226.6.6.6/32 PIM-SM:10.255.14.144, 239.2.0.0      RM
```

show mvpn c-multicast summary

```
user@host> show mvpn c-multicast summary
MVPN Summary:
Instance: VPN-A
  C-multicast IPv4 route count: 1
Instance: VPN-B
  C-multicast IPv4 route count: 2
```

show mvpn c-multicast extensive

```
user@host> show mvpn c-multicast extensive
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
  C-mcast IPv4 (S:G)          Ptnl          St
  192.168.195.78/32:225.5.5.5/32 PIM-SM:10.255.14.144, 239.1.1.1      RM
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-B
  C-mcast IPv4 (S:G)          Ptnl          St
  192.168.195.94/32:226.6.6.6/32 PIM-SM:10.255.14.144, 239.2.0.0      RM
```


show mvpn instance

Syntax	show mvpn instance <extensive summary> <instance <i>instance-name</i> >
Release Information	Command introduced in Junos OS Release 8.4.
Description	Display the multicast VPN routing instance information.
Options	extensive summary —(Optional) Display the specified level of output. instance <i>instance-name</i> —(Optional) Display statistics for the specified routing instance.
Required Privilege Level	view
List of Sample Output	show mvpn instance on page 1064 show mvpn instance on page 1064 show mvpn instance summary on page 1065 show mvpn instance extensive on page 1065 show mvpn instance summary (IPv6) on page 1065
Output Fields	Table 166 on page 1063 lists the output fields for the show mvpn instance command. Output fields are listed in the approximate order in which they appear.

Table 166: show mvpn instance Output Fields

Field Name	Field Description	Level of Output
MVPN instance	Name of the multicast VPN routing instance	extensive none
Instance	Name of the VPN routing instance.	summary extensive none
Provider tunnel	Provider tunnel attributes, <i>tunnel type:tunnel source, tunnel destination group</i> .	extensive none
Neighbor	Address, type of provider tunnel (I-P-tnl , inclusive provider tunnel and S-P-tnl , selective provider tunnel) and provider tunnel for each neighbor.	extensive none
C-mcast IPv4 (S:G)	Customer IPv4 router multicast address.	extensive none
C-mcast IPv6 (S:G)	Customer IPv6 router multicast address.	extensive none
Ptnl	Provider tunnel attributes, <i>tunnel type:tunnel source, tunnel destination group</i> .	extensive none
St	State: <ul style="list-style-type: none"> DS—Represents (S,G) and is created due to (*,G) RM—Remote VPN route learned from the remote PE router St display blank—SSM group join 	extensive none

Table 166: show mvpn instance Output Fields (*continued*)

Field Name	Field Description	Level of Output
Neighbor count	Number of neighbors associated with the multicast VPN routing instance.	summary
C-multicast IPv4 route count	Number of customer multicast IPv4 routes associated with the multicast VPN routing instance.	summary
C-multicast IPv6 route count	Number of customer multicast IPv6 routes associated with the multicast VPN routing instance.	summary

Sample Output

show mvpn instance

```

user@host> show mvpn instance
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
  Provider tunnel: I-P-tnl:PIM-SM:10.255.14.144, 239.1.1.1
  Neighbor
    10.255.14.160          I-P-tnl
    10.255.70.17          PIM-SM:10.255.14.160, 239.1.1.1
    10.255.70.17          PIM-SM:10.255.70.17, 239.1.1.1
  C-mcast IPv4 (S:G)      Ptnl          St
    192.168.195.78/32:225.5.5.5/32 PIM-SM:10.255.14.144, 239.1.1.1      RM
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-B
  Provider tunnel: I-P-tnl:PIM-SM:10.255.14.144, 239.2.0.0
  Neighbor
    10.255.14.160          I-P-tnl
    10.255.70.17          PIM-SM:10.255.14.160, 239.2.0.0
    10.255.70.17          PIM-SM:10.255.70.17, 239.2.0.0
  C-mcast IPv4 (S:G)      Ptnl          St
    192.168.195.94/32:226.6.6.6/32 PIM-SM:10.255.14.144, 239.2.0.0      RM

```

show mvpn instance

```

user@host> show mvpn instance
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance : vpn-1
  MVPN Mode : SPT-ONLY
  Provider tunnel: I-P-tnl:LDP-P2MP:10.255.72.162, lsp-id 16777217

```

Neighbor	I-P-tnl
10.255.72.160	LDP-P2MP:10.255.72.160, lsp-id 16777217
10.255.72.166	RSVP-TE P2MP:10.255.72.166,
13054,10.255.72.166	
10.255.72.168	

Sample Output

show mvpn instance summary

```
user@host> show mvpn instance summary
MVPN Summary:
Instance: VPN-A
  Neighbor count: 2
  C-multicast IPv4 route count: 1
Instance: VPN-B
  Neighbor count: 4
  C-multicast IPv4 route count: 2
```

Sample Output

show mvpn instance extensive

```
user@host> show mvpn instance extensive
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)      RM -- remote VPN route
Instance: VPN-A
  Provider tunnel: I-P-tnl:PIM-SM:10.255.14.144, 239.1.1.1
  Neighbor
    10.255.14.160      I-P-tnl      PIM-SM:10.255.14.160, 239.1.1.1
    10.255.70.17       PIM-SM:10.255.70.17, 239.1.1.1
  C-mcast IPv4 (S:G)      Ptnl      St
    192.168.195.78/32:225.5.5.5/32 PIM-SM:10.255.14.144, 239.1.1.1      RM
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)      RM -- remote VPN route
Instance: VPN-B
  Provider tunnel: I-P-tnl:PIM-SM:10.255.14.144, 239.2.0.0
  Neighbor
    10.255.14.160      I-P-tnl      PIM-SM:10.255.14.160, 239.2.0.0
    10.255.70.17       PIM-SM:10.255.70.17, 239.2.0.0
  C-mcast IPv4 (S:G)      Ptnl      St
    192.168.195.94/32:226.6.6.6/32 PIM-SM:10.255.14.144, 239.2.0.0      RM
```

show mvpn instance summary (IPv6)

```
user@host> show mvpn instance summary
MVPN Summary:
Instance: VPN-A
  C-multicast IPv6 route count: 2
Instance: VPN-B
  C-multicast IPv6 route count: 2
```


show mvpn neighbor

Syntax	show mvpn neighbor <extensive summary> <inet inet6> <instance <i>instance-name</i> neighbor-address <i>address</i> > <logical-system <i>logical-system-name</i> >
Release Information	Command introduced in Junos OS Release 8.4.
Description	Display multicast VPN neighbor information.
Options	<p>extensive summary—(Optional) Display the specified level of output for all multicast VPN neighbors.</p> <p>inet inet6—(Optional) Display IPv4 or IPv6 information for all multicast VPN neighbors.</p> <p>instance <i>instance-name</i> neighbor-address <i>address</i>—(Optional) Display multicast VPN neighbor information for the specified instance or the specified neighbor.</p> <p>logical-system <i>logical-system-name</i>—(Optional) Display multicast VPN neighbor information for the specified logical system.</p>
Required Privilege Level	view
List of Sample Output	show mvpn neighbor on page 1068 show mvpn neighbor extensive on page 1068 show mvpn neighbor extensive on page 1068 show mvpn neighbor instance-name on page 1069 show mvpn neighbor neighbor-address on page 1069 show mvpn neighbor neighbor-address summary on page 1069 show mvpn neighbor neighbor-address extensive on page 1070 show mvpn neighbor neighbor-address instance-name on page 1070
Output Fields	Table 167 on page 1067 lists the output fields for the show mvpn neighbor command. Output fields are listed in the approximate order in which they appear.

Table 167: show mvpn neighbor Output Fields

Field Name	Field Description	Level of Output
MVPN instance	Name of the multicast VPN routing instance	extensive none
Instance	Name of the VPN routing instance.	summary extensive none
Neighbor	Address, type of provider tunnel (I-P-tnl, inclusive provider tunnel and S-P-tnl, selective provider tunnel) and provider tunnel for each neighbor.	extensive none
Provider tunnel	Provider tunnel attributes, <i>tunnel type:tunnel source, tunnel destination group</i> .	extensive none

Sample Output

show mvpn neighbor

```
user@host> show mvpn neighbor
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
Neighbor                               I-P-tnl
10.255.14.160                          PIM-SM:10.255.14.160, 239.1.1.1
10.255.70.17                          PIM-SM:10.255.70.17, 239.1.1.1
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-B
Neighbor                               I-P-tnl
10.255.14.160                          PIM-SM:10.255.14.160, 239.2.0.0
10.255.70.17                          PIM-SM:10.255.70.17, 239.2.0.0
```

Sample Output

show mvpn neighbor extensive

```
user@host> show mvpn neighbor extensive
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
Neighbor                               I-P-tnl
10.255.14.160                          PIM-SM:10.255.14.160, 239.1.1.1
10.255.70.17                          PIM-SM:10.255.70.17, 239.1.1.1
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-B
Neighbor                               I-P-tnl
10.255.14.160                          PIM-SM:10.255.14.160, 239.2.0.0
10.255.70.17                          PIM-SM:10.255.70.17, 239.2.0.0
```

show mvpn neighbor extensive

```
user@host> show mvpn neighbor extensive
```

```

MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: mvpn-a
Neighbor                               I-P-tnl
10.255.72.45
10.255.72.50                           LDP P2MP:10.255.72.50, lsp-id 1

```

Sample Output

show mvpn neighbor instance-name

```

user@host> show mvpn neighbor instance-name VPN-A
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
Neighbor                               I-P-tnl
10.255.14.160                           PIM-SM:10.255.14.160, 239.1.1.1
10.255.70.17                            PIM-SM:10.255.70.17, 239.1.1.1

```

Sample Output

show mvpn neighbor neighbor-address

```

user@host> show mvpn neighbor neighbor-address 10.255.14.160
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
Neighbor                               I-P-tnl
10.255.14.160                           PIM-SM:10.255.14.160, 239.1.1.1
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-B
Neighbor                               I-P-tnl
10.255.14.160                           PIM-SM:10.255.14.160, 239.2.0.0

```

Sample Output

show mvpn neighbor neighbor-address summary

```

user@host> show mvpn neighbor neighbor-address 10.255.70.17 summary

```

MVPN Summary:
Instance: VPN-A
Instance: VPN-B

Sample Output

show mvpn neighbor neighbor-address extensive

```
user@host> show mvpn neighbor neighbor-address 10.255.70.17 extensive
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
Neighbor                             I-P-tnl
10.255.70.17                         PIM-SM:10.255.70.17, 239.1.1.1
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-B
Neighbor                             I-P-tnl
10.255.70.17                         PIM-SM:10.255.70.17, 239.2.0.0
```

Sample Output

show mvpn neighbor neighbor-address instance-name

```
user@host> show mvpn neighbor neighbor-address 10.255.70.17 instance-name VPN-A
MVPN instance:

Legend for provider tunnel
I-P-tnl -- inclusive provider tunnel S-P-tnl -- selective provider tunnel

Legend for c-multicast routes properties (Pr)
DS -- derived from (*, c-g)          RM -- remote VPN route
Instance: VPN-A
Neighbor                             I-P-tnl
10.255.70.17                         PIM-SM:10.255.70.17, 239.1.1.1
```


show vpls connections

Syntax	<pre>show vpls connections <brief extensive> <down up up-down> <history> <instance <i>instance-name</i> local-site <i>local-site-name</i> remote-site <i>remote-site-name</i>> <instance-history> <logical-system (all <i>logical-system-name</i>)> <status> <summary></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>instance-history option introduced in Junos OS Release 12.3R2.</p>
Description	(T Series and M Series routers, except for the M160 router) Display virtual private LAN service (VPLS) connection information.
Options	<p>none—Display information about all VPLS connections for all routing instances.</p> <p>brief extensive—(Optional) Display the specified level of output.</p> <p>down up up-down—(Optional) Display nonoperational, operational, or both types of connections.</p> <p>history—(Optional) Display information about connection history.</p> <p>instance <i>instance-name</i>—(Optional) Display the VPLS connections for the specified routing instance only.</p> <p>instance-history—(Optional) Display information about connection history for a particular instance.</p> <p>local-site <i>local-site-name</i>—(Optional) Display the VPLS connections for the specified local site name or ID only.</p> <p>remote-site <i>remote-site-name</i>—(Optional) Display the VPLS connections for the specified remote site name or ID only. Label block size information is always shown as 0 when using this option.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p> <p>status—(Optional) Display information about the connection and interface status.</p> <p>summary—(Optional) Display summary of all VPLS connections information.</p>
Required Privilege Level	view
List of Sample Output	<p>show vpls connections on page 1077</p> <p>show vpls connections extensive (Static VPLS Neighbors) on page 1079</p>

Output Fields Table 168 on page 1072 lists the output fields for the **show vpls connections** command. Output fields are listed in the approximate order in which they appear.

Table 168: show vpls connections Output Fields

Field Name	Field Description
Instance	Name of the VPLS instance.
Local site	Name of the local site.
VPLS-id	Identifier for the VPLS site.
Number of local interfaces	Number of interfaces configured for the local site.
Number of local interfaces up	Number of interfaces configured for the local site that are currently up.
IRB interface present	Indicates whether or not an integrated routing and bridging (IRB) interface is present (yes or no).
Intf	<p>List of all of the interfaces configured for the local site. The types of interfaces can include VPLS virtual loopback tunnel interfaces and label-switched interfaces. Any interface that supports VPLS could be listed here.</p> <p>Virtual loopback tunnel interfaces are displayed using the <i>vt-fpc/pic/port.nnnnn</i> format. Label-switched interfaces are displayed using the <i>lsi.nnnnn</i> format. In both cases, <i>nnnnn</i> is a dynamically generated virtual port used to transport and receive packets from other provider edge (PE) routers in the VPLS domain.</p> <p>Each interface might include the following information:</p> <ul style="list-style-type: none"> • Identification as a VPLS interface • Name of the associated VPLS routing instance • Local site number • Remote site number • VPLS neighbor address • VPLS identifier
Interface flags	<p>Flag associated with the interface. Can include the following:</p> <ul style="list-style-type: none"> • VC-Down—The virtual circuit associated with this interface is down.
Label-base	First label in a block of labels. A remote PE router uses this first label when sending traffic toward the advertising PE router.
Offset	Displays the VPLS Edge (VE) block offset in the Layer 2 VPN NLRI. The VE block offset is used to identify a label block from which a particular label value is selected to setup a pseduowire for a remote site. The block offset value itself indicates the starting VE ID that maps to the label base contained in the VPLS NLRI advertisement.

Table 168: show vpls connections Output Fields (*continued*)

Field Name	Field Description
Size	Label block size. A configurable value that represents the number of label blocks required to cover all the pseudowires for the remote peer. Acceptable configuration values are: 2 , 4 , 8 and 16 . The default value is 2 . A value of 0 will be displayed when using the remote-site option.
Range	Label block range.
Preference	Preference value advertised for a VPLS site. When multiple PE routers are assigned the same VE ID for multihoming, you might need to specify that a particular PE router acts as the designated forwarder by configuring the site preference value. The site preference indicates the degree of preference for a particular customer site. The site preference is one of the tie-breaking criteria used in a designated forwarder election.
status-vector	Bit vector advertising the state of local PE-CE circuits to remote PE routers. A bit value of 0 indicates that the local circuit and LSP tunnel to the remote PE router are up, whereas a value of 1 indicates either one or both are down.
connection-site	Name of the connection site.
Neighbor	IP address and VPLS identifier for the VPLS neighbor.
Type	Type of connection: loc (local) or rmt (remote).

Table 168: show vpls connections Output Fields (*continued*)

Field Name	Field Description
St	

Table 168: show vpls connections Output Fields (*continued*)

Field Name	Field Description
	<p>Status of the VPLS connection (corresponds with Legend for Connection Status):</p> <ul style="list-style-type: none"> • EI—The local VPLS interface is configured with an encapsulation that is not supported. • EM—The encapsulation type received on this VPLS connection from the neighbor does not match the local VPLS connection interface encapsulation type. • VC-Dn—The virtual circuit is currently down. • CM—The two routers do not agree on a control word, which causes a control word mismatch. • CN—The virtual circuit is not provisioned properly. • OR—The label associated with the virtual circuit is out of range. • OL—No advertisement has been received for this virtual circuit from the neighbor. There is no outgoing label available for use by this virtual circuit. • LD—All of the CE-facing interfaces to the local site are down. Therefore, the connection to the local site is signaled as down to the other PE routers. No pseudowires can be established. • RD—All the interfaces to the remote neighbor are down. Therefore, the remote site has been signaled as down to the other PE routers. No pseudowires can be established. • LN—The local site has lost path selection to the remote site and therefore no pseudowires can be established from this local site. • RN—The remote site has lost path selection to a local site or other remote site and therefore no pseudowires are established to this remote site. <p>In a multihoming configuration, one multihomed PE site displays the state LN, and the other multihomed PE site displays the state RN in the following circumstances:</p> <ul style="list-style-type: none"> • The multihomed links are both configured to be the backup site. • The two multihomed PE routers have the same site ID, but have a peering relationship with a route reflector (RR) that has a different site ID. <ul style="list-style-type: none"> • XX—The VPLS connection is down for an unknown reason. This is a programming error. • MM—The MTU for the local site and the remote site do not match. • BK—The router is using a backup connection. • PF—Profile parse failure. • RS—The remote site is in a standby state. • NC—The interface encapsulation is not configured as an appropriate CCC, TCC, or VPLS encapsulation. • WE—The encapsulation configured for the interface does not match the encapsulation configured for the associated connection within the VPLS routing instance. • NP—The router detects that interface hardware is not present. The hardware might be offline, a PIC might not be of the desired type, or the interface might be configured in a different routing instance. • ->—Only the outbound connection is up. • <--—Only the inbound connection is up. • Up—The VPLS connection is operational.

Table 168: show vpls connections Output Fields (*continued*)

Field Name	Field Description
	<ul style="list-style-type: none"> • Dn—The VPLS connection is down. • CF—The router cannot find enough bandwidth to the remote router to satisfy the VPLS connection bandwidth requirement. • SC—The local site identifier matches the remote site identifier. No pseudowire can be established between these two sites. You should configure different values for the local and remote site identifiers. • LM—The local site identifier is not the minimum designated, meaning it is not the lowest. There is another local site with a lower site identifier. Pseudowires are not being established to this local site, and the associated local site identifier is not being used to distribute VPLS label blocks. However, this is not an error state. Traffic continues to be forwarded to the PE router interfaces connected to the local sites when the local sites are in this state. • RM—The remote site identifier is not the minimum designated, meaning it is not the lowest. There is another remote site connected to the same PE router which has lower site identifier. The PE router cannot establish a pseudowire to this remote site and the associated remote site identifier cannot be used to distribute VPLS label blocks. However, this is not an error state. Traffic can continue to be forwarded to the PE router interface connected to this remote site when the remote site is in this state. • IL—The incoming packets for the VPLS connection have no MPLS label. • MI—The configured mesh group identifier is in use by another system in the network. • ST—The router has switched to a standby connection. • PB—Profile busy. • SN—The VPLS neighbor is static.
Time last up	Time connection was last in the Up condition.
# Up trans	Number of transitions from Down to Up condition.
Status	Status of the (local or remote circuit) local interface: <ul style="list-style-type: none"> • Up—Operational • Dn—Down • NP—Not present • DS—Disabled • WE—Wrong encapsulation • UN—Uninitialized
Encapsulation	Type of encapsulation: VPLS .
Remote PE	Address of the remote provider edge router.
Negotiated control-word	Whether a control word has been negotiated: Yes or No .
Incoming label	Name of the incoming label.
Outgoing label	Name of the outgoing label.

Table 168: show vpls connections Output Fields (*continued*)

Field Name	Field Description
Negotiated PW status TLV	Indicates whether or not the pseudowire status TLV has been negotiated for the VPLS connection.
Local interface	Provides the following information about the local interface configured for the VPLS neighbor: <ul style="list-style-type: none"> • Name of the local interface • Status—Interface status (Up or Down) • Encapsulation—Interface encapsulation (for example, ETHERNET) • Description—Includes the VPLS instance name, the VPLS neighbor address, and the VPLS identifier
Time	Date and time of VPLS connection event.
Event	Type of event.
Interface/Lbl/PE	Interface, label, or PE router.
Connection History	Each entry can include the date, time, year, and the connection event. Connection events include any of a variety of events related to VPLS connections, such as route changes, label updates, and interfaces going down or coming up.

Sample Output

show vpls connections

```
user@host> show vpls connections
```

```
Layer-2 VPN connections:
```

```
Legend for connection status (St)
```

```

EI -- encapsulation invalid      NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch    WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down  NP -- interface hardware not present
CM -- control-word mismatch    -< -- only outbound connection is up
CN -- circuit not provisioned  >- -- only inbound connection is up
OR -- out of range             Up -- operational
OL -- no outgoing label        Dn -- down
LD -- local site signaled down CF -- call admission control failure
RD -- remote site signaled down SC -- local and remote site ID collision
LN -- local site not designated LM -- local site ID not minimum designated
RN -- remote site not designated RM -- remote site ID not minimum designated
XX -- unn connection status    IL -- no incoming label
MM -- MTU mismatch             MI -- Mesh-Group ID not availble
BK -- Backup connection        ST -- Standby connection
PF -- Profile parse failure     PB -- Profile busy

```

```
Legend for interface status
```

```

Up -- operational
Dn -- down

```

```
Instance: vpls-1
```

```
Local site: 1 (11)
```

```
Number of local interfaces: 1
```

```

Number of local interfaces up: 1
IRB interface present: no
lt-1/3/0.10496
vt-1/3/0.1048588    1      Intf - vpls vpls-1 local site 11 remote site 1

vt-1/2/0.1048591    2      Intf - vpls vpls-1 local site 11 remote site 2

vt-1/2/0.1048585    3      Intf - vpls vpls-1 local site 11 remote site 3

vt-1/2/0.1048587    4      Intf - vpls vpls-1 local site 11 remote site 4

vt-1/2/0.1048589    5      Intf - vpls vpls-1 local site 11 remote site 5

vt-1/3/0.1048586    6      Intf - vpls vpls-1 local site 11 remote site 6

vt-1/3/0.1048590    7      Intf - vpls vpls-1 local site 11 remote site 7

vt-1/3/0.1048584    8      Intf - vpls vpls-1 local site 11 remote site 8

Label-base      Offset      Size      Range      Preference
+ 800256         1         16        16         100
Timer Values:
  Startup wait time: 120 seconds
  New site wait-time: 20 seconds
  Collision detect time: 30 seconds
  Reclaim wait time: 748 milliseconds
connection-site      Type      St      Time last up      # Up trans
1                    rmt      Up      Apr 28 13:28:24 2009      2
  Remote PE: 124.1.2.1, Negotiated control-word: No
  Incoming label: 800256, Outgoing label: 800026
  Local interface: vt-1/3/0.1048588, Status: Up, Encapsulation: VPLS
  Description: Intf - vpls vpls-1 local site 11 remote site 1
Connection History:
  Apr 28 13:28:24 2009 status update timer
  Apr 28 13:28:24 2009 PE route down
  Apr 28 13:24:27 2009 status update timer
  Apr 28 13:24:27 2009 loc intf up      vt-1/3/0.1048588
  Apr 28 13:24:27 2009 PE route changed
  Apr 28 13:24:27 2009 Out lbl Update      800026
  Apr 28 13:24:27 2009 In lbl Update      800256
  Apr 28 13:24:27 2009 loc intf down
2                    rmt      Up      Apr 28 13:28:24 2009      2
  Remote PE: 124.1.7.1, Negotiated control-word: No
  Incoming label: 800257, Outgoing label: 800034
  Local interface: vt-1/2/0.1048591, Status: Up, Encapsulation: VPLS
  Description: Intf - vpls vpls-1 local site 11 remote site 2
Connection History:
  Apr 28 13:28:24 2009 status update timer
  Apr 28 13:28:24 2009 PE route down
  Apr 28 13:24:28 2009 status update timer
  Apr 28 13:24:28 2009 loc intf up      vt-1/2/0.1048591
  Apr 28 13:24:28 2009 PE route changed
  Apr 28 13:24:28 2009 Out lbl Update      800034
  Apr 28 13:24:28 2009 In lbl Update      800257
  Apr 28 13:24:28 2009 loc intf down
3                    rmt      Up      Apr 28 13:28:24 2009      2
  Remote PE: 124.1.4.1, Negotiated control-word: No
  Incoming label: 800258, Outgoing label: 800026
  Local interface: vt-1/2/0.1048585, Status: Up, Encapsulation: VPLS
  Description: Intf - vpls vpls-1 local site 11 remote site 3
Connection History:

```



```

Apr 28 13:28:24 2009 status update timer
Apr 28 13:28:24 2009 PE route down
Apr 28 13:24:26 2009 status update timer
Apr 28 13:24:26 2009 loc intf up vt-1/2/0.1048585
Apr 28 13:24:26 2009 PE route changed
Apr 28 13:24:26 2009 Out lbl Update 800026
Apr 28 13:24:26 2009 In lbl Update 800258
Apr 28 13:24:26 2009 loc intf down
4 rmt Up Apr 28 13:28:24 2009 2
Remote PE: 124.1.6.1, Negotiated control-word: No
Incoming label: 800259, Outgoing label: 800026
Local interface: vt-1/2/0.1048587, Status: Up, Encapsulation: VPLS
Description: Intf - vpls vpls-1 local site 11 remote site 4
Connection History:
Apr 28 13:28:24 2009 status update timer
Apr 28 13:28:24 2009 PE route down
Apr 28 13:24:27 2009 status update timer
Apr 28 13:24:27 2009 loc intf up vt-1/2/0.1048587
Apr 28 13:24:27 2009 PE route changed
Apr 28 13:24:27 2009 Out lbl Update 800026
Apr 28 13:24:27 2009 In lbl Update 800259
Apr 28 13:24:27 2009 loc intf down
5 rmt Up Apr 28 13:28:24 2009 2
Remote PE: 124.1.3.1, Negotiated control-word: No
Incoming label: 800260, Outgoing label: 800034
Local interface: vt-1/2/0.1048589, Status: Up, Encapsulation: VPLS
Description: Intf - vpls vpls-1 local site 11 remote site 5
Connection History:
Apr 28 13:28:24 2009 status update timer
Apr 28 13:28:24 2009 PE route down
Apr 28 13:24:28 2009 status update timer
Apr 28 13:24:28 2009 loc intf up vt-1/2/0.1048589
Apr 28 13:24:28 2009 PE route changed
Apr 28 13:24:28 2009 Out lbl Update 800034
Apr 28 13:24:27 2009 In lbl Update 800260
Apr 28 13:24:27 2009 loc intf down

```

show vpls connections extensive (Static VPLS Neighbors)

```

user@host> show vpls connections extensive instance red
Layer-2 VPN connections:

```

Legend for connection status (St)

EI -- encapsulation invalid	NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch	WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down	NP -- interface hardware not present
CM -- control-word mismatch	-> -- only outbound connection is up
CN -- circuit not provisioned	<- -- only inbound connection is up
OR -- out of range	Up -- operational
OL -- no outgoing label	Dn -- down
LD -- local site signaled down	CF -- call admission control failure
RD -- remote site signaled down	SC -- local and remote site ID collision
LN -- local site not designated	LM -- local site ID not minimum designated
RN -- remote site not designated	RM -- remote site ID not minimum designated
XX -- unn connection status	IL -- no incoming label
MM -- MTU mismatch	MI -- Mesh-Group ID not availble
BK -- Backup connection	ST -- Standby connection
PF -- Profile parse failure	PB -- Profile busy
RS -- remote site standby	SN -- Static Neighbor

Legend for interface status

```

Up -- operational
Dn -- down

Instance: static
VPLS-id: 1
  Number of local interfaces: 1
  Number of local interfaces up: 1
  ge-0/0/5.0
  lsi.1049344          Intf - vpls static neighbor 10.255.114.3 vpls-id
1
Neighbor              Type St      Time last up      # Up trans
10.255.114.3(vpls-id 1)(SN) rmt Up    Mar  4 08:48:41 2010      1
  Remote PE: 10.255.114.3, Negotiated control-word: No
  Incoming label: 29696, Outgoing label: 29697
  Negotiated PW status TLV: No
  Local interface: lsi.1049344, Status: Up, Encapsulation: ETHERNET
  Description: Intf - vpls static neighbor 10.255.114.3 vpls-id 1
Connection History:
  Mar  4 08:48:41 2010  status update timer
  Mar  4 08:48:41 2010  PE route changed
  Mar  4 08:48:41 2010  Out lbl Update                      29697
  Mar  4 08:48:41 2010  In lbl Update                        29696
  Mar  4 08:48:41 2010  loc intf up                          lsi.1049344

```

```

user@PE1> show vpls connections extensive (Multihoming with FEC 129)
Layer-2 VPN connections:

```

Legend for connection status (St)

EI -- encapsulation invalid	NC -- interface encapsulation not CCC/TCC/VPLS
EM -- encapsulation mismatch	WE -- interface and instance encaps not same
VC-Dn -- Virtual circuit down	NP -- interface hardware not present
CM -- control-word mismatch	-> -- only outbound connection is up
CN -- circuit not provisioned	<- -- only inbound connection is up
OR -- out of range	Up -- operational
OL -- no outgoing label	Dn -- down
LD -- local site signaled down	CF -- call admission control failure
RD -- remote site signaled down	SC -- local and remote site ID collision
LN -- local site not designated	LM -- local site ID not minimum designated
RN -- remote site not designated	RM -- remote site ID not minimum designated
XX -- unknown connection status	IL -- no incoming label
MM -- MTU mismatch	MI -- Mesh-Group ID not available
BK -- Backup connection	ST -- Standby connection
PF -- Profile parse failure	PB -- Profile busy
RS -- remote site standby	SN -- Static Neighbor
LB -- Local site not best-site	RB -- Remote site not best-site
VM -- VLAN ID mismatch	

Legend for interface status

```

Up -- operational
Dn -- down

```

```

Instance: green
L2vpn-id: 100:100
Local-id: 1.1.1.2
  Number of local interfaces: 2
  Number of local interfaces up: 2
  ge-0/3/1.0
  ge-0/3/3.0
  lsi.101711873          Intf - vpls green local-id 1.1.1.2 remote-id
1.1.1.4 neighbor 1.1.1.4
  Remote-id              Type St      Time last up      # Up trans

```

```

1.1.1.4          rmt  Up    Jan 31 13:49:52 2012          1
Remote PE: 1.1.1.4, Negotiated control-word: No
Incoming label: 262146, Outgoing label: 262146
Local interface: lsi.101711873, Status: Up, Encapsulation: ETHERNET
Description: Intf - vpls green local-id 1.1.1.2 remote-id 1.1.1.4 neighbor
1.1.1.4
Connection History:
  Jan 31 13:49:52 2012  status update timer
  Jan 31 13:49:52 2012  PE route changed
  Jan 31 13:49:52 2012  Out lbl Update                262146
  Jan 31 13:49:52 2012  In lbl Update                  262146
  Jan 31 13:49:52 2012  loc intf up                    lsi.101711873
Multi-home:
Local-site      Id    Pref  State
test            1    100   Up
Number of interfaces: 1
Number of interfaces up: 1
ge-0/3/1.0
Received multi-homing advertisements:
Remote-PE      Pref  flag  Description
1.1.1.4        100   0x0

```

show vpls flood event-queue

Syntax	show vpls flood event-queue
Release Information	Command introduced in Junos OS Release 8.0.
Description	Display the pending events in the VPLS flood queue.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show vpls flood event-queue on page 1082
Output Fields	Table 169 on page 1082 lists the output fields for the show vpls flood event-queue command. Output fields are listed in the approximate order in which they appear.

Table 169: show vpls flood event-queue Output Fields

Field Name	Field Description
Current Pending Event	Provides information on the current event in the VPLS flood event queue.
Name	Name of the event.
Owner Name	Name of the interface associated with the flood event.
Pending Op	Pending operation for the event.
Last Error	Name of the last error encountered.
Number of Retries	Number of attempts made to update the event queue.
Pending Event List	List of the events awaiting processing.
Event Name	Name of the event.
Pending Op	Pending operation for the event.
Event Identifier	Name of the interface associated with the flood event.

Sample Output

show vpls flood event-queue

```

user@host> show vpls flood event-queue
Current Pending Event
  Name:          Flood Nexthop
  Owner Name: ge-4/3/0.0
  Pending Op: ADD

```

```
Last Error:ENOMEM
Number of Retries:3
Pending Event List:
Event Name      Pending Op      Event Identifier
Flood Nexthop   ADD             ge-4/3/0.0
Flood Route     ADD             ge-4/3/0.0
```

show vpls flood instance

Syntax	show vpls flood instance <brief detail extensive> <instance-name> <logical-system <i>logical-system-name</i> >
Release Information	Command introduced in Junos OS Release 8.0.
Description	Display VPLS information related to the flood process.
Options	<p>none—Display VPLS information related to the flood process for all routing instances.</p> <p>brief detail extensive—(Optional) Display the specified level of output.</p> <p>instance-name—(Optional) Display VPLS information related to the flood process for the specified routing instance.</p> <p>logical-system <i>logical-system-name</i>—(Optional) Display VPLS information related to the flood process for the specified logical system.</p>
Required Privilege Level	view
List of Sample Output	show vpls flood instance on page 1085 show vpls flood instance logical-system-name on page 1085 show vpls flood instance detail on page 1085
Output Fields	Table 170 on page 1084 lists the output fields for the show vpls flood instance command. Output fields are listed in the approximate order in which they appear.

Table 170: show vpls flood instance Output Fields

Field Name	Field Description
Logical system	Name of the logical system.
Name	Name of the VPLS routing instance.
CEs	Number of CE routers connected to the VPLS instance.
VEs	Number of VE routers connected to the VPLS instance.
Flood routes	List of all flood routes associated with the VPLS instance.
Prefix	Prefix for the route.
Type	Type of route.
Owner	VPLS routing instance or interface associated with the route.
Nhype	Next-hop type. For example, flood for a flood route.

Table 170: show vpls flood instance Output Fields (*continued*)

Field Name	Field Description
Nhindex	Next-hop index number for the route.

Sample Output

show vpls flood instance

```

user@host> show vpls flood instance

Logical system: __juniper_ls1__
Name: green
CEs: 1
VEs: 1
Flood Routes:
  Prefix   Type      Owner      NhType      NhIndex
  default  ALL_CE_FLOOD green      flood       383
  0x47/16  CE_FLOOD  fe-1/2/1.0 flood       388

```

show vpls flood instance logical-system-name

```

user@host: __juniper_ls1__> show vpls flood instance juniper_ls1

Logical system: __juniper_ls1__
Name: green
CEs: 1
VEs: 1
Flood Routes:
  Prefix   Type      Owner      NhType      NhIndex
  default  ALL_CE_FLOOD green      flood       383
  0x47/16  CE_FLOOD  fe-1/2/1.0 flood       388

```

show vpls flood instance detail

```

user@host: __juniper_ls1__> show vpls flood instance detail

Logical system: __juniper_ls1__
Name: green
CEs: 1
VEs: 1
Flood Routes:
  Prefix   Type      Owner      NhType      NhIndex
  default  ALL_CE_FLOOD green      flood       383
  0x47/16  CE_FLOOD  fe-1/2/1.0 flood       388

```

show vpls flood route

Syntax	show vpls flood route (all-ce-flood instance-name <i>instance-name</i> <logical-system-name <i>logical-system-name</i> > ce-flood interface <i>interface-name</i>)
Release Information	Command introduced in Junos OS Release 8.0.
Description	Display VPLS route information related to the flood process for either the specified routing instance or the specified interface.
Options	<p>all-ce-flood—Display the flood next-hop route for all customer edge routers for traffic coming from the core of the network.</p> <p>ce-flood interface <i>interface-name</i>—Display the flood next-hop route for traffic coming from the specified customer edge interface.</p> <p>instance-name <i>instance-name</i>—Display the flood routes for the specified instance.</p> <p>logical-system-name <i>logical-system-name</i>—(Optional) Specify the logical system whose flood routes you want to display. You can only specify the default logical system name for VPLS. The default logical system name is __juniper_ls1__ (the name must be entered in the command with the underscore characters).</p>
Required Privilege Level	view
List of Sample Output	show vpls flood route all-ce-flood on page 1087 show vpls flood route ce-flood on page 1087
Output Fields	Table 171 on page 1086 lists the output for the show vpls flood route command. Output fields are listed in the approximate order in which they appear.

Table 171: show vpls flood route Output Fields

Field Name	Field Description
Flood route prefix	Prefix for the flood route.
Flood route type	Type of flood route (either CE_FLOOD or ALL_CE_FLOOD).
Flood route owner	VPLS routing instance or interface associated with the flood route.
Nexthop type	Next-hop type. For example, flood for a flood route.
Nexthop index	Next-hop index number for the route.
Interfaces flooding to	Interfaces to which VPLS routes are being flooded.
Name	Name of the interface.

Table 171: show vpls flood route Output Fields (*continued*)

Field Name	Field Description
Type	Type of VPLS router (CE or VE).
Nh type	Next-hop type.
Index	Index number for the flood route.

Sample Output

show vpls flood route all-ce-flood

```
user@host:~juniper_ls1~> show vpls flood route all-ce-flood logical-system-name
~juniper_ls1_instance-name green
```

```
Flood route prefix: default
Flood route type: ALL_CE_FLOOD
Flood route owner: green
Nexthop type: flood
Nexthop index: 383
  Interfaces Flooding to:
    Name      Type      NhType      Index
    fe-1/2/1.0 CE
```

show vpls flood route ce-flood

```
user@host:~juniper_ls1~> show vpls flood route ce-flood interface fe-1/2/1.0
```

```
Flood route prefix: 0x47/16
Flood route type: CE_FLOOD
Flood route owner: fe-1/2/1.0
Nexthop type: flood
Nexthop index: 388
  Interfaces Flooding to:
    Name      Type      NhType      Index
    lsi.49152 VE      indr      262142
```

show vpls mac-table

Syntax	<pre>show vpls mac-table <brief detail extensive summary> <bridge-domain <i>bridge-domain-name</i>> <instance <i>instance-name</i>> <interface <i>interface-name</i>> <logical-system (all <i>logical-system-name</i>)> <mac-address> <vlan-id <i>vlan-id-number</i>></pre>
Release Information	Command introduced in Junos OS Release 8.5.
Description	(MX960 routers only) Display learned VPLS MAC address information.
Options	<p>none—Display all learned VPLS MAC address information.</p> <p>brief detail extensive summary—(Optional) Display the specified level of output.</p> <p>bridge-domain <i>bridge-domain-name</i>—(Optional) Display learned VPLS MAC addresses for the specified bridge domain.</p> <p>instance <i>instance-name</i>—(Optional) Display learned VPLS MAC addresses for the specified instance.</p> <p>interface <i>interface-name</i>—(Optional) Display learned VPLS MAC addresses for the specified instance.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Display learned VPLS MAC addresses for all logical systems or for the specified logical system.</p> <p>mac-address—(Optional) Display the specified learned VPLS MAC address information..</p> <p>vlan-id <i>vlan-id-number</i>—(Optional) Display learned VPLS MAC addresses for the specified VLAN.</p>
Required Privilege Level	view
List of Sample Output	<p>show vpls mac-table on page 1089</p> <p>show vpls mac-table (with VXLAN enabled) on page 1090</p> <p>show vpls mac-table count on page 1090</p> <p>show vpls mac-table detail on page 1091</p> <p>show vpls mac-table extensive on page 1091</p>
Output Fields	<p>Table 172 on page 1088 describes the output fields for the show bridge mac-table command. Output fields are listed in the approximate order in which they appear.</p>

Table 172: show vpls mac-table Output fields

Field Name	Field Description
Routing instance	Name of the routing instance.

Table 172: show vpls mac-table Output fields (*continued*)

Field Name	Field Description
Bridging domain	Name of the bridging domain.
MAC address	MAC address or addresses learned on a logical interface.
MAC flags	Status of MAC address learning properties for each interface: <ul style="list-style-type: none"> • S—Static MAC address configured. • D—Dynamic MAC address learned. • SE—MAC accounting is enabled. • NM—Nonconfigured MAC.
Logical interface	Name of the logical interface.
MAC count	Number of MAC addresses learned on a specific routing instance or interface.
Learning interface	Logical interface or logical Label Switched Interface (LSI) the address is learned on.
Learn VLAN ID/VLAN	VLAN ID of the routing instance or bridge domain in which the MAC address was learned.
VXLAN ID/VXLAN	VXLAN Network Identifier (VNI)
Layer 2 flags	Debugging flags signifying that the MAC address is present in various lists.
Epoch	Spanning Tree Protocol epoch number identifying when the MAC address was learned. Used for debugging.
Sequence number	Sequence number assigned to this MAC address. Used for debugging.
Learning mask	Mask of Packet Forwarding Engines where this MAC address was learned. Used for debugging.
IPC generation	Creation time of the logical interface when this MAC address was learned. Used for debugging.

Sample Output

show vpls mac-table

```

user@host> show vpls mac-table
MAC flags (S -static MAC, D -dynamic MAC,
          SE -Statistics enabled, NM -Non configured MAC)

Routing instance : vpls_ldp1
VLAN : 223
  MAC          MAC      Logical
  address      flags     interface
  00:90:69:9c:1c:5d  D      ge-0/2/5.400

MAC flags (S -static MAC, D -dynamic MAC,
          SE -Statistics enabled, NM -Non configured MAC)

Routing instance : vpls_red

```

```

VLAN : 401
MAC          MAC      Logical
address      flags    interface
00:00:aa:12:12:12 D      lsi.1051138
00:05:85:74:9f:f0 D      lsi.1051138

```

show vpls mac-table (with VXLAN enabled)

```

user@host> show vpls mac-table
MAC flags (S -static MAC, D -dynamic MAC, L -locally learned
          SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

Routing instance : vpls_4site:1000
Bridging domain : __vpls_4site:1000__, VLAN : 4094,4093
VXLAN: Id : 300, Multicast group: 226.1.1.3
MAC          MAC      Logical
address      flags    interface
00:01:01:00:01:f4 D,SE    ge-4/2/0.1000
00:02:01:33:01:f4 D,SE    lsi.1052004
00:03:00:32:01:f4 D,SE    lsi.1048840
00:04:00:14:01:f4 D,SE    lsi.1052005
00:02:01:33:02:f7 D,SE    vtep.1052010
00:04:00:14:02:f7 D,SE    vtep.1052011

```

show vpls mac-table count

```

user@host> show vpls mac-table count
0 MAC address learned in routing instance __juniper_private1__

MAC address count per interface within routing instance:
Logical interface      MAC count
lc-0/0/0.32769         0
lc-0/1/0.32769         0
lc-0/2/0.32769         0
lc-2/0/0.32769         0
lc-0/3/0.32769         0
lc-2/1/0.32769         0
lc-9/0/0.32769         0
lc-11/0/0.32769        0
lc-2/2/0.32769         0
lc-9/1/0.32769         0
lc-11/1/0.32769        0
lc-2/3/0.32769         0
lc-9/2/0.32769         0
lc-11/2/0.32769        0
lc-11/3/0.32769        0
lc-9/3/0.32769         0

MAC address count per learn VLAN within routing instance:
Learn VLAN ID          MAC count
0                       0

1 MAC address learned in routing instance vpls_ldp1

MAC address count per interface within routing instance:
Logical interface      MAC count
lsi.1051137            0
ge-0/2/5.400           1

MAC address count per learn VLAN within routing instance:
Learn VLAN ID          MAC count

```

	0	1
1 MAC address learned in routing instance vpls_red		
MAC address count per interface within routing instance:		
Logical interface	MAC count	
ge-0/2/5.300	1	
MAC address count per learn VLAN within routing instance:		
Learn VLAN ID	MAC count	
0	1	

show vpls mac-table detail

```

user@host> show vpls mac-table detail
MAC address: 00:90:69:9c:1c:5d
Routing instance: vpls_ldp1
Learning interface: ge-0/2/5.400
Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
Epoch: 0                               Sequence number: 1
Learning mask: 0x1                       IPC generation: 0

MAC address: 00:90:69:9c:1c:5d
Routing instance: vpls_red
Learning interface: ge-0/2/5.300
Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
Epoch: 0                               Sequence number: 1
Learning mask: 0x1                       IPC generation: 0

```

show vpls mac-table extensive

```

user@host> show vpls mac-table extensive
MAC address: 00:00:aa:12:12:12
Routing instance: vpls_ldp1
Learning interface: lsi.1051137
Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
Epoch: 0                               Sequence number: 1
Learning mask: 0x1                       IPC generation: 0

MAC address: 00:05:85:74:9f:f0
Routing instance: vpls_ldp1
Learning interface: lsi.1051137
Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
Epoch: 0                               Sequence number: 1
Learning mask: 0x1                       IPC generation: 0

MAC address: 00:90:69:9c:1c:5d
Routing instance: vpls_ldp1
Learning interface: ge-0/2/5.400
Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
Epoch: 0                               Sequence number: 1
Learning mask: 0x1                       IPC generation: 0

MAC address: 00:00:aa:12:12:12
Routing instance: vpls_red
Learning interface: lsi.1051138
Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
Epoch: 0                               Sequence number: 0
Learning mask: 0x1                       IPC generation: 0

MAC address: 00:05:85:74:9f:f0

```

```
Routing instance: vpls_red
Learning interface: lsi.1051138
Layer 2 flags: in_ifd, in_ifl, in_vlan, kernel
Epoch: 0                               Sequence number: 0
Learning mask: 0x1                       IPC generation: 0
```

show vpls statistics

Syntax	show vpls statistics <instance <i>instance-name</i> > <logical-system (all <i>logical-system-name</i>)>
Release Information	Command introduced before Junos OS Release 7.4.
Description	(T Series and M Series routers, except for the M160 router) Display virtual private LAN service (VPLS) statistics.
Options	<p>none—Display VPLS statistics for all routing instances.</p> <p>instance <i>instance-name</i>—(Optional) Display VPLS statistics for a specific VPLS routing instance only.</p> <p>logical-system (all <i>logical-system-name</i>)—(Optional) Perform this operation on all logical systems or on a particular logical system.</p>
Required Privilege Level	view
List of Sample Output	show vpls statistics on page 1094 show vpls statistics instance on page 1094
Output Fields	Table 173 on page 1093 lists the output fields for the show vpls statistics command. Output fields are listed in the approximate order in which they appear.

Table 173: show vpls statistics Output Fields

Field Name	Field Description
Instance	Name of the VPLS instance.
Local interface	Name of the local VPLS virtual loopback tunnel interface, <i>vt-fpc/pic/port.nnnnn</i> , where <i>nnnnn</i> is a dynamically generated virtual port used to transport and receive packets from other provider edge (PE) routers in the VPLS domain.
Index	Number associated with the next hop.
Remote provider edge router	Address of the remote provider edge router.
Multicast packets	Number of multicast packets received.
Multicast bytes	Number of multicast bytes received.
Flood packets	Number of VPLS flood packets received.
Flood bytes	Number of VPLS flood bytes received.

Table 173: show vpls statistics Output Fields (*continued*)

Field Name	Field Description
Current MAC count	Number of MAC addresses learned by the interface and the configured maximum limit on the number of MAC addresses that can be learned.

Sample Output

show vpls statistics

```

user@host> show vpls statistics

VPLS statistics:

Instance: green

  Local interface: fe-2/2/1.0, Index: 69
    Multicast packets:      1
    Multicast bytes   :    60
    Flooded packets   :    18
    Flooded bytes    :   2556
    Current MAC count:      1

  Local interface: lt-0/3/0.2, Index: 72
    Multicast packets:      3
    Multicast bytes   :   153
    Flooded packets   :      1
    Flooded bytes    :    51
    Current MAC count:      1

  Local interface: lsi.32769, Index: 75
    Current MAC count:      0

  Local interface: lsi.32771, Index: 77
  Remote PE: 10.255.14.222
    Current MAC count:      2

Instance: red

  Local interface: vt-0/3/0.32768, Index: 74
    Multicast packets:      0
    Multicast bytes   :      0
    Flooded packets   :      0
    Flooded bytes    :      0
    Current MAC count:      0

  Local interface: vt-0/3/0.32770, Index: 76
    Multicast packets:      0
    Multicast bytes   :      0
    Flooded packets   :      0
    Flooded bytes    :      0
    Current MAC count:      0

```

show vpls statistics instance

```

user@host> show vpls statistics instance red

```


Layer-2 VPN Statistics:

Instance: red

Local interface: vt-3/2/0.32768, Index: 73

Remote provider edge router: 10.255.17.35

Multicast packets: 0

Multicast bytes : 0

Flood packets : 0

Flood bytes : 0

Current MAC count: 1 (Limit 20)

PART 4

Troubleshooting

- [Interface Diagnostics on page 1099](#)

Interface Diagnostics

- [Interface Diagnostics on page 1099](#)

Interface Diagnostics

You can use two diagnostic tools to test the physical layer connections of interfaces: loopback testing and bit error rate test (BERT) testing. Loopback testing enables you to verify the connectivity of a circuit. BERT testing enables you to identify poor signal quality on a circuit. This section contains the following topics:

- [Configuring Loopback Testing on page 1099](#)
- [Interface Diagnostics on page 1101](#)

Configuring Loopback Testing

Loopback testing allows you to verify the connectivity of a circuit. You can configure any of the following interfaces to execute a loopback test: Aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, E1, E3, NxDS0, serial, SONET/SDH, T1, and T3.

The physical path of a network data circuit usually consists of segments interconnected by devices that repeat and regenerate the transmission signal. The transmit path on one device connects to the receive path on the next device. If a circuit fault occurs in the form of a line break or a signal corruption, you can isolate the problem by using a loopback test. Loopback tests allow you to isolate segments of the circuit and test them separately.

To do this, configure a *line loopback* on one of the routers. Instead of transmitting the signal toward the far-end device, the line loopback sends the signal back to the originating router. If the originating router receives back its own data link layer packets, you have verified that the problem is beyond the originating router. Next, configure a line loopback farther away from the local router. If this originating router does not receive its own data link layer packets, you can assume the problem is on one of the segments between the local router and the remote router's interface card. In this case, the next troubleshooting step is to configure a line loopback closer to the local router to find the source of the problem.

There are several types of loopback testing supported by the Junos OS, as follows:

- DCE local—Loops packets back on the local DCE.
- DCE remote—Loops packets back on the remote DCE.

- **Local**—Useful for troubleshooting physical PIC errors. Configuring local loopback on an interface allows transmission of packets to the channel service unit (CSU) and then to the circuit toward the far-end device. The interface receives its own transmission, which includes data and timing information, on the local router's PIC. The data received from the CSU is ignored. To test a local loopback, issue the **show interfaces *interface-name*** command. If PPP keepalives transmitted on the interface are received by the PIC, the **Device Flags** field contains the output **Loop-Detected**.
- **Payload**—Useful for troubleshooting the physical circuit problems between the local router and the remote router. A payload loopback loops data only (without clocking information) on the remote router's PIC. With payload loopback, overhead is recalculated.
- **Remote**—Useful for troubleshooting the physical circuit problems between the local router and the remote router. A remote loopback loops packets, including both data and timing information, back on the remote router's interface card. A router at one end of the circuit initiates a remote loopback toward its remote partner. When you configure a remote loopback, the packets received from the physical circuit and CSU are received by the interface. Those packets are then retransmitted by the PIC back toward the CSU and the circuit. This loopback tests all the intermediate transmission segments.

Table 174 on page 1100 shows the loopback modes supported on the various interface types.

Table 174: Loopback Modes by Interface Type

Interface	Loopback Modes	Usage Guidelines
Aggregated Ethernet, Fast Ethernet, Gigabit Ethernet	Local	<i>Configuring Ethernet Loopback Capability</i>
Circuit Emulation E1	Local and remote	<i>Configuring E1 Loopback Capability</i>
Circuit Emulation T1	Local and remote	<i>Configuring T1 Loopback Capability</i>
E1 and E3	Local and remote	<i>Configuring E1 Loopback Capability and Configuring E3 Loopback Capability</i>
NxDS0	Payload	<i>Configuring Channelized E1 IQ and IQE Interfaces, Configuring T1 and NxDS0 Interfaces, Configuring Channelized OC12/STM4 IQ and IQE Interfaces (SONET Mode), Configuring Channelized STM1 IQ and IQE Interfaces, and Configuring Channelized T3 IQ Interfaces</i>
Serial (V.35 and X.21)	Local and remote	<i>Configuring Serial Loopback Capability</i>
Serial (EIA-530)	DCE local, DCE remote, local, and remote	<i>Configuring Serial Loopback Capability</i>
SONET/SDH	Local and remote	<i>Configuring SONET/SDH Loopback Capability</i>

Table 174: Loopback Modes by Interface Type (*continued*)

Interface	Loopback Modes	Usage Guidelines
T1 and T3	Local, payload, and remote	<i>Configuring T1 Loopback Capability</i> and <i>Configuring T3 Loopback Capability</i> <i>See also Configuring the T1 Remote Loopback Response</i>

To configure loopback testing, include the **loopback** statement:

loopback mode;

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* aggregated-ether-options]
- [edit interfaces *interface-name* ds0-options]
- [edit interfaces *interface-name* e1-options]
- [edit interfaces *interface-name* e3-options]
- [edit interfaces *interface-name* fastether-options]
- [edit interfaces *interface-name* gigether-options]
- [edit interfaces *interface-name* serial-options]
- [edit interfaces *interface-name* sonet-options]
- [edit interfaces *interface-name* t1-options]
- [edit interfaces *interface-name* t3-options]

Interface Diagnostics

BERT allows you to troubleshoot problems by checking the quality of links. You can configure any of the following interfaces to execute a BERT when the interface receives a request to run this test: E1, E3, T1, T3; the channelized DS3, OC3, OC12, and STM1 interfaces; and the channelized DS3 IQ, E1 IQ, and OC12 IQ interfaces.

A BERT test requires a line loop to be in place on either the transmission devices or the far-end router. The local router generates a known bit pattern and sends it out the transmit path. The received pattern is then verified against the sent pattern. The higher the bit error rate of the received pattern, the worse the noise is on the physical circuit. As you move the position of the line loop increasingly downstream toward the far-end router, you can isolate the troubled portion of the link.

To configure BERT, you must configure the duration of the test, the bit pattern to send on the transmit path, and the error rate to monitor when the inbound pattern is received.

To configure the duration of the test, the pattern to send in the bit stream, and the error rate to include in the bit stream, include the **bert-period**, **bert-algorithm**, and **bert-error-rate** statements, respectively, at the [edit interfaces *interface-name* *interface-type*-options] hierarchy level:

```
[edit interfaces interface-name interface-type-options]
bert-algorithm algorithm;
bert-error-rate rate;
bert-period seconds;
```

By default, the BERT period is 10 seconds. You can configure the BERT period to last from 1 through 239 seconds on some PICs and from 1 through 240 seconds on other PICs.

rate is the bit error rate. This can be an integer from 0 through 7, which corresponds to a bit error rate from 10^{-0} (1 error per bit) to 10^{-7} (1 error per 10 million bits).

algorithm is the pattern to send in the bit stream. For a list of supported algorithms, enter a ? after the **bert-algorithm** statement; for example:

```
[edit interfaces t1-0/0/0 t1-options]
user@host# set bert-algorithm ?
Possible completions:
pseudo-2e11-o152      Pattern is 2^11 - 1 (per 0.152 standard)
pseudo-2e15-o151      Pattern is 2^15 - 1 (per 0.152 standard)
pseudo-2e20-o151      Pattern is 2^20 - 1 (per 0.151 standard)
pseudo-2e20-o153      Pattern is 2^20 - 1 (per 0.153 standard)
...
```

For specific hierarchy information, see the individual interface types.



NOTE: The 4-port E1 PIC supports only the following algorithms:

pseudo-2e11-o152	Pattern is 2^11 - 1 (per 0.152 standard)
pseudo-2e15-o151	Pattern is 2^15 - 1 (per 0.151 standard)
pseudo-2e20-o151	Pattern is 2^20 - 1 (per 0.151 standard)
pseudo-2e23-o151	Pattern is 2^23 (per 0.151 standard)

When you issue the help command from the CLI, all BERT algorithm options are displayed, regardless of the PIC type, and no commit check is available. Unsupported patterns for a PIC type can be viewed in system log messages.



NOTE: The 12-port T1/E1 Circuit Emulation (CE) PIC supports only the following algorithms:

```
all-ones-repeating    Repeating one bits
all-zeros-repeating   Repeating zero bits
alternating-double-ones-zeros Alternating pairs of ones and zeros
alternating-ones-zeros Alternating ones and zeros
pseudo-2e11-o152     Pattern is 2^11 - 1 (per 0.152 standard)
pseudo-2e15-o151     Pattern is 2^15 - 1 (per 0.151 standard)
pseudo-2e20-o151     Pattern is 2^20 - 1 (per 0.151 standard)
pseudo-2e7           Pattern is 2^7 - 1
pseudo-2e9-o153      Pattern is 2^9 - 1 (per 0.153 standard)
repeating-1-in-4      1 bit in 4 is set
repeating-1-in-8      1 bit in 8 is set
repeating-3-in-24     3 bits in 24 are set
```

When you issue the help command from the CLI, all BERT algorithm options are displayed, regardless of the PIC type, and no commit check is available. Unsupported patterns for a PIC type can be viewed in system log messages.



NOTE: The IQE PICs support only the following algorithms:

```
all-ones-repeating    Repeating one bits
all-zeros-repeating   Repeating zero bits
alternating-double-ones-zeros Alternating pairs of ones and zeros
alternating-ones-zeros Alternating ones and zeros
pseudo-2e9-o153       Pattern is 2^9 - 1 (per 0.153 (511 type) standard)
pseudo-2e11-o152      Pattern is 2^11 - 1 (per 0.152 and 0.153 (2047 type)
standards)
pseudo-2e15-o151      Pattern is 2^15 - 1 (per 0.151 standard)
pseudo-2e20-o151      Pattern is 2^20 - 1 (per 0.151 standard)
pseudo-2e20-o153      Pattern is 2^20 - 1 (per 0.153 standard)
pseudo-2e23-o151      Pattern is 2^23 - 1 (per 0.151 standard)
repeating-1-in-4       1 bit in 4 is set
repeating-1-in-8       1 bit in 8 is set
repeating-3-in-24      3 bits in 24 are set
```

When you issue the help command from the CLI, all BERT algorithm options are displayed, regardless of the PIC type, and no commit check is available. Unsupported patterns for a PIC type can be viewed in system log messages.



NOTE: BERT is supported on the PDH interfaces of the Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP and the DS3/E3 MIC. The following BERT algorithms are supported:

all-ones-repeating	Repeating one bits
all-zeros-repeating	Repeating zero bits
alternating-double-ones-zeros	Alternating pairs of ones and zeros
alternating-ones-zeros	Alternating ones and zeros
repeating-1-in-4	1 bit in 4 is set
repeating-1-in-8	1 bit in 8 is set
repeating-3-in-24	3 bits in 24 are set
pseudo-2e9-o153	Pattern is $2^9 - 1$ (per 0.153 standard)
pseudo-2e11-o152	Pattern is $2^{11} - 1$ (per 0.152 standard)
pseudo-2e15-o151	Pattern is $2^{15} - 1$ (per 0.151 standard)
pseudo-2e20-o151	Pattern is $2^{20} - 1$ (per 0.151 standard)
pseudo-2e20-o153	Pattern is $2^{20} - 1$ (per 0.153 standard)
pseudo-2e23-o151	Pattern is $2^{23} - 1$ (per 0.151 standard)

Table 175 on page 1104 shows the BERT capabilities for various interface types.

Table 175: BERT Capabilities by Interface Type

Interface	T1 BERT	T3 BERT	Comments
12-port T1/E1 Circuit Emulation	Yes (ports 0–11)		<ul style="list-style-type: none"> Limited algorithms
4-port Channelized OC3/STM1 Circuit Emulation	Yes (port 0–3)		<ul style="list-style-type: none"> Limited algorithms
E1 or T1	Yes (port 0–3)	Yes (port 0–3)	<ul style="list-style-type: none"> Single port at a time Limited algorithms
E3 or T3	Yes (port 0–3)	Yes (port 0–3)	<ul style="list-style-type: none"> Single port at a time
Channelized OC12	N/A	Yes (channel 0–11)	<ul style="list-style-type: none"> Single channel at a time Limited algorithms No bit count
Channelized STM1	Yes (channel 0–62)	N/A	<ul style="list-style-type: none"> Multiple channels Only one algorithm No error insert No bit count
Channelized T3 and Multichannel T3	Yes (channel 0–27)	Yes (port 0–3 on channel 0)	<ul style="list-style-type: none"> Multiple ports and channels Limited algorithms for T1 No error insert for T1 No bit count for T1

These limitations do not apply to channelized IQ interfaces. For information about BERT capabilities on channelized IQ interfaces, see *Channelized IQ and IQE Interfaces Properties*.

Starting and Stopping a BERT Test

Before you can start the BERT test, you must disable the interface. To do this, include the **disable** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
disable;
```

After you configure the BERT properties and commit the configuration, begin the test by issuing the **test interface *interface-name* *interface-type*-bert-start** operational mode command:

```
user@host> test interface interface-name interface-type-bert-start
```

The test runs for the duration you specify with the **bert-period** statement. If you wish to terminate the test sooner, issue the **test interface *interface-name* *interface-type*-bert-stop** command:

```
user@host> test interface interface-name interface-type-bert-stop
```

For example:

```
user@host> test interface t3-1/2/0 t3-bert-start
user@host> test interface t3-1/2/0 t3-bert-stop
```

To view the results of the BERT test, issue the **show interfaces extensive | find BERT** command:

```
user@host> show interfaces interface-name extensive | find BERT
```

For more information about running and evaluating the results of the BERT procedure, see the [CLI Explorer](#).



NOTE: To exchange BERT patterns between a local router and a remote router, include the **loopback remote** statement in the interface configuration at the remote end of the link. From the local router, issue the **test interface** command.

Example: Configuring Bit Error Rate Testing

Configure a BERT test on a T3 interface. In this example, the run duration lasts for 120 seconds. The configured error rate is 0, which corresponds to a bit error rate of 10^{-0} (1 error per bit). The configured bit pattern of **all-ones-repeating** means that every bit the interface sends is a set to a value of 1.

```
[edit interfaces]
t3-1/2/0 {
  t3-options {
    bert algorithm all-ones-repeating;
    bert-error-rate 0;
    bert-period 120;
```

```
}  
}
```

PART 5

Index

- [Index on page 1109](#)

Index

Symbols

#, comments in configuration statements.....	xxvii
(), in syntax descriptions.....	xxvii
128-bit IPv6 address.....	39
32-bit IPv4 address.....	37
< >, in syntax descriptions.....	xxvii
[], in configuration statements.....	xxvii
{ }, in configuration statements.....	xxvii
(pipe), in syntax descriptions.....	xxvii

A

accounting statement.....	105
usage guidelines.....	25
address statement.....	102
usage guidelines.....	36, 44
advertisements, displaying	783
aggregate statement	
hierarchical policer.....	104
AMT	
protocol, displaying.....	401
state	
clearing.....	363
statistics	
clearing.....	362
statistics, displaying.....	398
tunnel	
clearing.....	363
tunnel, displaying.....	403
ANCP	
clearing statistics.....	262
CoS state	
displaying ANCP.....	268
local access loop	
displaying.....	284
neighbor	
clearing.....	260
neighbor summary	
displaying.....	291
neighbors	
displaying.....	271
request OAM per interface.....	266

request OAM per neighbor.....	267
statistics	
displaying.....	279
subscriber	
clearing.....	264
subscriber summary	
displaying.....	293
subscribers	
displaying.....	284
summary	
displaying.....	289
arp option	
policers.....	11
arp statement.....	106
AS paths	
distribution of, displaying.....	765
domain information, displaying.....	769
matching regular expressions, displaying.....	790
summary of, displaying.....	771
ASN	
BGP community routes, displaying.....	798
autonomous system number See ASN	
autonomous system paths See AS paths	

B

backup PE groups	
multicast, displaying.....	457
bandwidth-limit statement	
hierarchical policer.....	107
BERT	
configuring interface diagnostics.....	1101
bert-algorithm statement	
usage guidelines.....	1101
bert-error-rate statement	
usage guidelines.....	1101
bert-period statement	
usage guidelines.....	1101
best routes, displaying.....	792
BFD	
sessions	
clearing.....	297
displaying.....	298
BGP	
community ASN, displaying routes.....	798
community name, displaying routes.....	800
damping parameters	
clearing.....	308
displaying.....	347
damping routes, displaying.....	802

groups	
general information, displaying.....	315
traffic statistics, displaying.....	322
neighbors	
clearing connections.....	309
displaying.....	324
summary information, displaying.....	341
table	
clearing.....	311
BGP Monitoring Protocol	
displaying	
statistics.....	313
Bidirectional Forwarding Detection See BFD	
bit error rate test See BERT	
BOOTP	
accepting packets.....	60
bootstrap routers, displaying.....	507
borrower interface	
unnumbered Ethernet or demux.....	45
braces, in configuration statements.....	xxvii
brackets	
angle, in syntax descriptions.....	xxvii
square, in configuration statements.....	xxvii
bridge domain	
Layer 2, displaying.....	1013
broadcast statement.....	108
usage guidelines.....	36
bundle statement.....	109
burst-size-limit statement	
hierarchical policer.....	110
C	
cbr statement.....	111
clear (ospf ospf3) database command.....	686
clear (ospf ospf3) io-statistics command.....	689
clear (ospf ospf3) neighbor command.....	690
clear (ospf ospf3) overload command.....	692
clear (ospf ospf3) statistics command.....	693
clear amt tunnel command.....	362, 363
clear ancp neighbor command.....	260
clear ancp statistics command.....	262
clear ancp subscriber command.....	264
clear bfd adaptation command.....	296
clear bfd session command.....	297
clear bgp damping command.....	308
clear bgp neighbor command.....	309
clear bgp table command.....	311
clear bridge mac-table command.....	1010
clear error bpdv command.....	1011
clear error mac-rewrite command.....	1012
clear esis adjacency command.....	350
clear esis statistics command.....	351
clear firewall command.....	988
clear igmp membership command.....	364
clear igmp snooping membership command.....	367
clear igmp snooping statistics command.....	368
clear igmp statistics command.....	369
clear ipv6 neighbors command.....	590
clear ipv6 router-advertisement command.....	591
clear isis adjacency command.....	598
clear isis database command.....	600
clear isis overload command.....	602
clear isis statistics command.....	604
clear lldp neighbor command.....	656
clear lldp statistics command.....	657
clear mld membership command.....	371
clear mld statistics command.....	372
clear msdp cache command.....	373
clear msdp statistics command.....	374
clear multicast bandwidth-admission	
command.....	375
clear multicast forwarding-cache command.....	377
clear multicast scope command.....	378
clear multicast sessions command.....	379
clear multicast snooping statistics command.....	380
clear multicast statistics command.....	381
clear pgm negative-acknowledgments	
command.....	382
clear pgm source-path-messages command.....	383
clear pgm statistics command.....	384
clear pim join command.....	385
clear pim join-distribution command.....	387
clear pim register command.....	389
clear pim statistics command.....	395
clear rip general-statistics command.....	970
clear rip statistics command.....	971
clear ripng general-statistics command.....	980
clear ripng statistics command.....	981
clear vpls mac-address command.....	1038
clear vpls mac-table command.....	1039
comments, in configuration statements.....	xxvii
community ASN, displaying routes.....	798
community name, displaying routes.....	800
conventions	
text and syntax.....	xxvi
curly braces, in configuration statements.....	xxvii
customer support.....	xxviii
contacting JTAC.....	xxviii

D

damping parameters, BGP	
clearing.....	308
displaying.....	347
damping routes, BGP	
displaying.....	802
default router addresses.....	50
demux interfaces	
unnumbered.....	45
demux0 statement	
dynamic IP demux interface.....	112
destination statement	
tunnels.....	114
usage guidelines.....	36, 43
destination-class usage	
example configuration.....	29
destination-class-usage statement.....	115
usage guidelines.....	25
destination-profile	
usage guidelines.....	43
destination-profile statement.....	115
usage guidelines.....	36
DHCP	
accepting.....	60
documentation	
comments on.....	xxvii
donor interface	
unnumbered Ethernet or demux.....	45
Draft-rosen MVPNs	
data MDT cache, displaying.....	537
MDT tunnels	
displaying.....	539
DVMRP	
groups, displaying.....	498
interfaces, displaying.....	406
neighbors, displaying.....	408
prefixes, displaying.....	410
prunes, displaying active.....	412
Dynamic Host Configuration Protocol See DHCP	
dynamic overload bit, resetting for IS-IS.....	602
dynamic profiles statements	
dynamic-profiles.....	116
interfaces.....	139
unnumbered-address.....	183
dynamic subscribers	
interfaces statement.....	139
dynamic tunnel database, displaying.....	1041

dynamic-profiles	
interfaces statement.....	139
dynamic IP demux.....	139
dynamic-profiles statement.....	116

E

epd-threshold statement	
logical interface.....	123
ES-IS	
adjacencies	
clearing.....	350
displaying.....	352
interfaces, displaying.....	354
statistics	
clearing.....	351
displaying.....	356
Ethernet interfaces	
status information, displaying	
Fast Ethernet.....	219
Gigabit Ethernet.....	194, 235
unnumbered.....	45
preferred source address.....	156
eui-64 statement.....	124
usage guidelines.....	36
export route information, displaying.....	826

F

fail-filter statement.....	163
usage guidelines.....	52
family bridge	
VLAN ID list.....	189
VLAN IDs.....	189
family statement.....	127
dynamic profiles.....	125
usage guidelines.....	33
Fast Ethernet interfaces	
status information, displaying.....	219
fast-aps-switch statement.....	131
filter statement.....	132
usage guidelines.....	20
firewall	
filter version	
displaying.....	997
hierarchical-policer.....	134
statistics	
displaying.....	990
firewall filters	
applying.....	20
example configuration.....	23

log information, displaying.....	998	IGMP	
logical interfaces.....	20	group membership, displaying.....	414
policed packets, displaying.....	1005	interfaces, displaying.....	418
statistics		PIM-to-IGMP message translation information,	
clearing.....	988	displaying.....	470
displaying.....	1001	snooping (interface).....	422
flooding		snooping (membership).....	425
Layer 2 bridging, displaying.....	1015	snooping (statistics).....	429
font conventions.....	xxvi	statistics, displaying.....	432
forward-and-send-to-re statement.....	133	inet protocol family	
forward-only statement.....	133	interface addresses.....	37
forwarding table		inet6 protocol family	
interfaces, displaying.....	864	interface addresses.....	39
multicast information, displaying.....	474	input option	
multicast snooping information,		firewall filters.....	20
displaying.....	492	policers.....	11
route entries, displaying.....	850	input statement.....	135
		input-list statement.....	136
G		usage guidelines.....	20
Gigabit Ethernet interfaces		interface addresses	
status information, displaying.....	194, 235	logical interfaces.....	36
Gigabit Ethernet IQ PIC		preferred interface addresses.....	50, 51
traffic and MAC statistics.....	194	primary interface addresses.....	50
group option		interface groups.....	22
firewall filters.....	20	interface learning information	
groups		Layer 2, displaying.....	1034
BGP		interface-mode statement.....	137
general information, displaying.....	315	interfaces	
traffic statistics, displaying.....	322	configuration statements.....	76
DVMRP, displaying.....	498	firewall filters.....	20
IGMP membership, displaying.....	414	unit statement.....	173
MLD		interfaces statement.....	138
clearing.....	371	dynamic profiles.....	139
displaying.....	435	interfaces, displaying	
PIM		in the forwarding table.....	864
general information, displaying.....	512	Internet Protocol Control Protocol See IPCP	
usage information, displaying.....	498	invalid routes, displaying.....	889
		inverse-arp statement.....	143
H		IP addresses	
hidden routes, displaying.....	867	128-bit.....	39
hierarchical-policer statement.....	134	32-bit.....	37
hostnames		IPCP.....	42
IS-IS, displaying.....	634	unnumbered interfaces.....	44
		IP IGMP snooping	
I		membership	
if-exceeding statement		clearing.....	367
hierarchical policer.....	135	statistics	
		clearing.....	368

IP multicast	
announced sessions, displaying.....	489
backup PE groups, displaying.....	457
bandwidth admission	
clearing.....	375
flow map information, displaying.....	459
forwarding cache, clearing.....	377
forwarding table, displaying.....	474
forwarding-cache statistics	
displaying.....	461
interface information, displaying.....	463
network information, displaying.....	465
next-hop table, displaying.....	467
PIM-to-IGMP message translation information,	
displaying.....	470
PIM-to-MLD message translation information,	
displaying.....	472
RPF calculations, displaying.....	481
SAP announcements, displaying.....	585
scope, clearing.....	378
scoped information, displaying.....	485
sessions, clearing.....	379
statistics	
clearing.....	381
displaying.....	495
IP multicast snooping	
forwarding table, displaying.....	492
statistics	
clearing.....	380
displaying.....	487
IPCP.....	42
assigning PPP properties.....	43
configuring IP address.....	42
negotiating IP addresses.....	43
unnumbered interfaces.....	43
ipsec-sa statement.....	143
IPv4 Protocol family	
interface addresses.....	36
Same IP Address on Multiple	
Interfaces.....	38
on logical interfaces.....	33
IPv6.....	35
neighbor cache information	
clearing.....	590
displaying.....	592
router advertisements	
clearing.....	591
displaying.....	594
standards documents.....	35
transition.....	35
IPv6 Protocol family	
on logical interfaces.....	33
IS-IS	
adjacency database entries, clearing.....	598
authentication, displaying.....	610
backup coverage	
displaying.....	612
backup MPLS LSPs.....	614
backup paths	
SPF calculations.....	616
dynamic overload bit, resetting.....	602
hostname database, displaying.....	634
interfaces, displaying.....	636
link-state database entries	
clearing.....	600
displaying.....	622
neighbors, displaying.....	606
routes, displaying.....	643
SPF calculations, displaying.....	647
traffic statistics	
clearing.....	604
displaying.....	652
ISO Protocol family.....	33, 36
J	
join states, clearing PIM.....	385
join states, redistributing.....	387
K	
keep-address-and-control statement.....	144
usage guidelines.....	10
L	
Layer 2 BPDU loop prevention	
clearing errors.....	1011
Layer 2 bridging	
bridge domain, displaying.....	1013
flooding, displaying.....	1015
interface learning information,	
displaying.....	1034
learning information, displaying.....	1029, 1031
learning properties, displaying.....	1032
MAC address table, clearing.....	1010
MAC address table, displaying.....	1022
statistics, displaying.....	1026

Layer 2 circuits	
connections, displaying.....	1048
redundant pseudowires.....	1040
Layer 2 protocol tunneling	
clearing errors.....	1012
interface information, displaying.....	1036
Layer 2 VPNs	
connections, displaying.....	1055
VPLS connections, displaying.....	1071
VPLS statistics, displaying.....	1093
Layer 3 VPNs	
dynamic tunnel database, displaying.....	1041
PIM MDTs, displaying.....	533
learning information	
Layer 2, displaying.....	1029, 1031
learning properties	
Layer 2, displaying.....	1032
LLDP	
displaying.....	658
local information, displaying.....	661
neighbor	
clearing.....	656
neighbors, displaying.....	664
remote global statistics, displaying.....	669
statistics	
clearing.....	657
statistics, displaying.....	671
logical interface statements	
family.....	125
logical interfaces	
default router addresses.....	50
firewall filters.....	20
interface addresses.....	36, 50
policers.....	11
preferred interface addresses.....	50, 51
primary interface addresses.....	50, 51
primary router addresses.....	50
primary router interfaces.....	50
protocol families.....	33
protocol MTU.....	9
protocol redirect messages.....	10
unnumbered interfaces.....	44
logical systems	
configuration statements.....	92
logical-systems statement.....	144
loopback testing.....	1099
M	
MAC address table	
Layer 2, clearing address information.....	1010
Layer 2, displaying.....	1022
manuals	
comments on.....	xxvii
martians, displaying.....	889
MDT	
displaying information.....	533
MDT join TLV	
displaying advertisements received.....	537
MLD	
group membership	
clearing.....	371
displaying.....	435
interfaces, displaying.....	439
PIM-to-MLD message translation information,	
displaying.....	472
statistics	
clearing.....	372
displaying.....	443
mode statement.....	145
usage guidelines.....	56
MPLS	
labels, displaying routes.....	882
protocol family.....	33, 36
MSDP	
cache entries, clearing.....	373
general information, displaying.....	446
message source information, displaying.....	448
peer statistics	
clearing.....	374
displaying.....	453
source-active cache, displaying.....	450
mtu statement.....	146
logical interfaces	
usage guidelines.....	9
MTUs	
logical interfaces.....	9
protocol MTUs.....	9
multicast distribution trees See MDT	
multicast VPNs	
customer multicast routes, displaying.....	1061
neighbors, displaying.....	1067
routing instances, displaying.....	1063
multicast-only statement.....	148
usage guidelines.....	34
multipoint-destination statement.....	149

MVPN.....	537, 539
See also Draft-Rosen MVPNs	
MVRP	
configuration	
show.....	674, 676, 679, 680
dynamic-vlan-memberships	
show.....	678
statistics	
show.....	682

N

negotiate-address statement.....	150
usage guidelines.....	43
negotiating IP addresses	
IPCP.....	43
neighbor	
ANCP.....	260
LLDP.....	656
next hops	
multicast entries, displaying.....	467
resolution database, displaying.....	930
routes sent to, displaying.....	892
no-redirects statement.....	150
usage guidelines.....	10
no-translate-discard-eligible statement.....	171
no-translate-fecn-and-becn statement.....	172

O

oam-liveness statement.....	151
oam-period statement.....	152
Open Shortest Path First See OSPF	
OSPF	
backup coverage	
displaying.....	695
backup MPLS LSPs	
displaying.....	698
backup neighbor paths.....	700
backup paths	
SPF calculations.....	702
context identifier, displaying.....	709
interfaces, displaying.....	730
link-state database entries, displaying	
version 2.....	711
version 3.....	719
neighbors	
clearing connections.....	690
displaying.....	741
overview	
displaying.....	747

routing table entries, displaying.....	752
SPF calculations, displaying.....	738
statistics, general	
clearing.....	693
displaying.....	758
statistics, I/O	
clearing.....	689
displaying.....	736
output option	
firewall filters.....	20
policers.....	11
output statement.....	153
output-list statement.....	153
usage guidelines.....	20
overload bit, resetting for IS-IS.....	602

P

parentheses, in syntax descriptions.....	xxvii
PGM	
negative acknowledgments	
clearing.....	382
displaying.....	501
source path messages	
clearing.....	383
displaying.....	503
statistics	
clearing.....	384
displaying.....	504
PIM	
bootstrap routers, displaying.....	507
groups	
general information, displaying.....	512
usage information, displaying.....	498
interfaces	
displaying.....	509
join states, clearing.....	385
MDTs, displaying.....	533
neighbors, displaying.....	541
PIM-to-IGMP message translation information,	
displaying.....	470
PIM-to-MLD message translation information,	
displaying.....	472
prune states, clearing.....	385
redistributing join states.....	387
register	
clearing.....	389
RPF, displaying source state.....	569

RPs	
displaying.....	545
statistics	
clearing.....	395
displaying.....	572
point-to-point connections	
unnumbered Ethernet interfaces.....	44
policer	
interface.....	154
policer statement	
usage guidelines.....	11
policers	
applying.....	11
arp option.....	11
burst-size-limit	
statement.....	11
input option.....	11
logical interfaces.....	11
output option.....	11
policers, displaying.....	1005
post-service-filter statement.....	155
PPP properties, assigning	
IPCP.....	43
preferred interface addresses.....	50, 51
preferred statement.....	155
usage guidelines.....	37, 51
preferred-source-address statement.....	156
example.....	49
usage guidelines.....	46
premium statement	
hierarchical policer.....	157
primary interface addresses.....	50, 51
primary router addresses.....	50
primary router interfaces.....	50
primary statement	
address for interface	
usage guidelines.....	51
address on interface.....	158
interface for router	
usage guidelines.....	50
protocol families	
logical interfaces.....	33
protocol MTUs.....	9
unnumbered interfaces.....	44
protocol MTUs	
logical interfaces.....	9
protocol redirect messages.....	10
protocols statement.....	158
proxy statement.....	159
prune states, clearing PIM.....	385
prunes, DVMRP, displaying.....	412
pseudowires	
redundant switchover.....	1040
Q	
queue-length statement.....	160
R	
receive-options-packets statement.....	160
receive-ttl-exceeded statement.....	161
redirect messages.....	10
redundant pseudowires, switchover.....	1040
regular expressions	
AS paths, displaying matching routes.....	790
IP multicast scope	
clearing.....	378
IP multicast sessions	
clearing.....	379
displaying.....	489
target communities, displaying.....	829
remote statement.....	161
rendezvous points <i>See</i> RPs	
replication	
of BGP configuration.....	338
request ancp oam interface command.....	266
request ancp oam neighbor command.....	267
request l2circuit-switchover command.....	1040
request pim multicast-tunnel rebalance	
command.....	397
reverse path forwarding <i>See</i> RPF	
RIP	
general statistics	
clearing.....	970
displaying.....	972
neighbors	
displaying.....	974
statistics	
clearing.....	971
displaying.....	976
RIPng	
general statistics	
clearing.....	980
displaying.....	982
neighbors	
displaying.....	983
statistics	
clearing.....	981
displaying.....	985

route advertisements, displaying.....	783	summary statistics.....	947
route, displaying		terse information.....	965
next-hop.....	892	RPF	
router advertisements		calculations, displaying.....	481
IPv6		PIM source state, displaying.....	569
clearing.....	591	rpf-check statement.....	162, 163
displaying.....	594	usage guidelines.....	52
routers		RPs	
default addresses.....	50	displaying.....	545
primary addresses.....	50	rtvbr statement.....	164
primary interfaces.....	50		
routes, displaying		S	
active.....	773	sampling statement.....	165
active path.....	778	SAP session announcements, displaying.....	585
advertising protocol.....	783	service statement	
all.....	788	logical interfaces.....	166
AS paths		service-filter statement.....	167
distribution of.....	765	service-set statement.....	167
domain information.....	769	Session Announcement Protocol See SAP	
regular expressions, matching.....	790	shaping statement.....	168
summary of.....	771	shortest path first.....	647
best.....	792	show (ospf ospf3) backup coverage	
brief information.....	795	command.....	695
CCC.....	797	show (ospf ospf3) backup lsp.....	698
community ASN.....	798	show (ospf ospf3) backup neighbor.....	700
community name.....	800	show (ospf ospf3) backup spf.....	702
damping, BGP.....	802	show (ospf ospf3) interface command.....	730
detailed information.....	807	show (ospf ospf3) io-statistics command.....	736
extensive information.....	831	show (ospf ospf3) log command.....	738
flow validation.....	848	show (ospf ospf3) neighbor command.....	741
hidden.....	867	show (ospf ospf3) overview command.....	747
in a prefix range.....	918	show (ospf ospf3) route command.....	752
in a specific routing table.....	951	show (ospf ospf3) statistics command.....	758
in the forwarding table.....	850	show amt statistics command.....	398
inactive path.....	870	show amt summary command.....	401
inactive prefix.....	873	show amt tunnel command.....	403
instances.....	875	show ancp cos command.....	268
learned from a specific address.....	941	show ancp neighbor command.....	271
learned from a specific protocol.....	906	show ancp statistics command.....	279
learned from snooping.....	933	show ancp subscriber command.....	284
LSP.....	885	show ancp summary command.....	289
martian.....	889	show ancp summary neighbor command.....	291
matching the specified address.....	824	show ancp summary subscriber command.....	293
MPLS labels.....	882	show as-path command.....	765
next-hop resolution.....	930	show as-path domain command.....	769
not associated with a community.....	898	show as-path summary command.....	771
policy-based route export.....	826	show bfd session command.....	298
received through a neighbor.....	922	show bgp bmp command.....	313
sent to a specific interface.....	901	show bgp bgp command.....	315

show bgp group traffic-statistics command.....	322
show bgp neighbor command.....	324
show bgp replication command.....	338
show bgp summary command.....	341
show bridge domain command.....	1013
show bridge flood command.....	1015
show bridge mac-table command.....	1022
show bridge statistics command.....	1026
show command.....	674, 676, 679, 680
show dvmrp interfaces command.....	406
show dvmrp neighbors command.....	408
show dvmrp prefix command.....	410
show dvmrp prunes command.....	412
show dynamic-tunnels database command.....	1041
show esis adjacency command.....	352
show esis interface command.....	354
show esis statistics command.....	356
show firewall command.....	990
show firewall filter version command.....	997
show firewall log command.....	998
show firewall prefix-action-stats command.....	1001
show firewall templates-in-use command.....	1003
show igmp group command.....	414
show igmp interface command.....	418
show igmp snooping interface command.....	422
show igmp snooping membership command.....	425
show igmp snooping statistics command.....	429
show igmp statistics command.....	432
show interfaces (10-Gigabit Ethernet) command.....	194
show interfaces (Fast Ethernet) command.....	219
show interfaces (Gigabit Ethernet) command.....	235
show ipv6 neighbors command.....	592
show ipv6 router-advertisement command.....	594
show isis adjacency command.....	606
show isis authentication command.....	610
show isis backup coverage.....	612
show isis backup spf results.....	616
show isis context-identifier command.....	620
show isis database command.....	622
show isis hostname command.....	634
show isis interface command.....	636
show isis overview command.....	640
show isis route command.....	643
show isis spf command.....	647
show isis statistics command.....	652
show isis-backup label-switched-path command.....	614
show l2-learning global-information command.....	1029
show l2-learning global-mac-count command.....	1031
show l2-learning instance command.....	1032
show l2-learning interface command.....	1034
show l2circuit connections command.....	1048
show l2vpn connections command.....	1055
show lldp command.....	658
show lldp local-information command.....	661
show lldp neighbors command.....	664
show lldp remote-global-statistics command.....	669
show lldp statistics command.....	671
show mac-rewrite interface command.....	1036
show mld group command.....	435
show mld interface command.....	439
show mld statistics command.....	443
show msdp command.....	446
show msdp source command.....	448
show msdp source-active command.....	450
show msdp statistics command.....	453
show multicast backup-pe-groups command.....	457
show multicast flow-map command.....	459
show multicast forwarding-cache statistics command.....	461
show multicast interface command.....	463
show multicast mrinfo command.....	465
show multicast next-hops command.....	467
show multicast pim-to-igmp-proxy command.....	470
show multicast pim-to-mld-proxy command.....	472
show multicast route command.....	474
show multicast rpf command.....	481
show multicast scope command.....	485
show multicast sessions command.....	489
show multicast snooping next-hops command.....	487
show multicast snooping route command.....	492
show multicast statistics command.....	495
show multicast usage command.....	498
show mvpn c-multicast command.....	1061
show mvpn instance command.....	1063
show mvpn neighbor command.....	1067
show mvrp statistics command.....	682
show ospf context-identifier command.....	709
show ospf database command.....	711
show ospf3 database command.....	719

show pgm negative-acknowledgments command.....	501
show pgm source-path-messages command.....	503
show pgm statistics command.....	504
show pim bootstrap command.....	507
show pim interfaces command.....	509
show pim join command.....	512
show pim mdt command.....	533
show pim mdt data-mdt-joins command.....	537
show pim mdt data-mdt-limit command.....	539
show pim neighbors command.....	541
show pim rps command.....	545
show pim source command.....	569
show pim statistics command.....	572
show policer command.....	1005
show policy damping command.....	347
show rip general-statistics command.....	972
show rip neighbor command.....	974
show rip statistics command.....	976
show ripng general-statistics command.....	982
show ripng neighbor command.....	983
show ripng statistics command.....	985
show route active-path command.....	778
show route advertising-protocol command.....	783
show route all command.....	788
show route aspath-regex command.....	790
show route best command.....	792
show route brief command.....	795
show route ccc command.....	797
show route command.....	773
show route community command.....	798
show route community-name command.....	800
show route damping command.....	802
show route detail command.....	807
show route exact command.....	824
show route export command.....	826
show route export vrf-target command.....	829
show route extensive command.....	831
show route flow validation command.....	848
show route forwarding-table command.....	850
show route forwarding-table interface-name command.....	864
show route hidden command.....	867
show route inactive-path command.....	870
show route inactive-prefix command.....	873
show route instance command.....	875
show route label command.....	882
show route label-switched-path command.....	885
show route martians command.....	889
show route next-hop command.....	892
show route no-community command.....	898
show route output command.....	901
show route protocol command.....	906
show route range command.....	918
show route receive-protocol command.....	922
show route resolution command.....	930
show route snooping command.....	933
show route source-gateway command.....	941
show route summary command.....	947
show route table command.....	951
show route terse command.....	965
show sap listen command.....	585
show vpls connections command.....	1071
show vpls flood event-queue command.....	1082
show vpls flood instance command.....	1084
show vpls flood route command.....	1086
show vpls mac-table command.....	1088
show vpls statistics command.....	1093
snooping (interface)	
IGMP.....	422
snooping (membership)	
IGMP.....	425
snooping (statistics)	
IGMP.....	429
snooping routes, displaying.....	933
source gateway addresses, displaying.....	941
source-class usage	
example configuration.....	29
source-class-usage statement.....	169
usage guidelines.....	25
SPF calculations, displaying.....	738
static routes	
unnumbered Ethernet interfaces.....	47
static subscribers	
interfaces statement.....	139
statistics	
Layer 2 bridging, displaying.....	1026
LLDP.....	657
subscriber	
ANCP.....	264
subscriber interface statements	
demux0.....	112
family.....	125
interfaces.....	139
mode.....	145
rpf-check.....	162
unit.....	173
unnumbered-address.....	183

support, technical See technical support

syntax conventions.....xxvi

T

targeted-broadcast statement.....170

technical support

 contacting JTAC.....xxviii

test msdp command.....587

then statement

 hierarchical policer.....171

translate-discard-eligible statement.....171

translate-fecn-and-becn statement.....172

Trivial Network Protocol family

 interface addresses.....36

 on logical interfaces.....33

trunk interface.....137

trunk port.....137

 VLAN ID list.....189

U

unicast RPF.....52

 example configuration.....60, 65

 fail filters.....54, 60, 65

 loose mode.....56

 routing asymmetry.....54

 strict mode.....54

 VPNs.....59

 example configuration.....60

unicast-reverse-path statement

 usage guidelines.....65

unit statement.....176

 interfaces.....173

unnumbered interfaces

 demux.....45

 Ethernet.....45

 preferred source address.....156

 IPCP.....43

 point-to-point.....44

unnumbered-address statement

 demux interface.....182

 dynamic profiles.....183

 Ethernet.....185

 PPP.....186

 preferred source address

 usage guidelines.....46

 usage guidelines.....43, 45

unnumbered-interface statement

 usage guidelines.....43

V

vbr statement.....187

vci statement.....188

verification

 static route.....71

virtual private LAN service See VPLS

vlan-id statement

 interface in bridge domain.....189

vlan-id-list statement

 bridge domain.....189

VPLS

 connections, displaying.....1071

 statistics, displaying.....1093

vpls protocol family

 interface addresses.....36

 on logical interfaces.....33

VPN routing and forwarding See VRF

VPNs

 unicast RPF.....59

VRF

 route export, displaying.....829